Toward a Theory of Harms in the Internet Ecosystem

David D. Clark and kc claffy

August 26, 2019

Contents

1 What do we want our future Internet to be, and what are the barriers to achieving it?

It is time to reckon with the implications of 30 years of a laissez-faire approach to Internet regulation. The hands-off approach, largely led by U.S. policy decisions starting in the 1990s, has facilitated unprecedented innovation in information and communication technologies, and brought rich connectivity and services to most of the developed world, and much of the developing world. But with most of the world using it, the Internet is now a different ecosystem, incorporating inevitable elements of the human condition that motivated the invention of regulations in the first place. Policy approaches that were reasonable in the past now come at the cost of real harms and threats to individuals, organizations, and society at large. To those without access, the now undeniable harm is marginalization with respect to the growing fraction of civil, commercial, and social activity that occurs online, sometimes primarily online. To those with access, the list of potential harms expands considerably: cybercrime, online bullying, and device addiction top a long list of negative impacts of the otherwise magical concept of ubiquitous and continuous connectivity to information and interaction. The most profound harms articulated about the new digital information ecosystem affect individuals whether or not they engage with any particular online service: lack of trust and privacy; political polarization; death of trusted journalism; consolidation and centralization of power and capital; slowed economic growth, productivity, and innovation; national security vulnerabilities; and corruption of democratic political and electoral processes. There seems to be no limit to the range of society's ills that someone will argue are amplified or enabled by the Internet.

Regulation of the Internet ecosystem is officially a hot topic – in the United States it has already become a catalyzing campaign issue for the 2020 presidential election. But there is still lively deliberation in traditional mainstream media, the blogosphere and podcast media, academia, and regulatory agencies themselves regarding exactly *what* to regulate and *how*.

We are not policy experts, but have read enough Internet policy literature to reluctantly conclude that the public policy community lacks comprehensive theories about harms in the Internet ecosystem, and often also lacks supporting empirical data to substantiate arguments for new legislation or to guide regulatory intervention under existing frameworks. We attempt a contribution toward filling this gap, in part informed by surveying recent work in this area, and in part by our decades of research and thinking about Internet architecture, technology, operations, engineering, economics, and policy.

We have previously undertaken an aspirational approach to considering the future of the Internet – surveying societal goals for Internet infrastructure and services. We identified and classified sixteen high-level aspirations, as a first step toward being able to measure progress toward such goals.[1] This exercise illustrated two facts that are already well-known to policy makers. First, the path from aspiration to accomplishment is often not obvious. Second, aspirations are often in conflict – embedding tradeoffs implicit in every Internet policy debate we observe today. Not everyone shares each aspiration, but such a list provides a starting point for more rigorous debate about how, or if, to develop policy tools to pursue a given set of aspirations.

In this work, we take the inverse approach: cataloging barriers or impediments to achieving aspirations. Some barriers are intrinsic – the speed of light will constrain synchronization capabilities across a wide area network, whether for gaming, music performance, or telesurgery. Some barriers are economic – universal gigabit access might be desirable, but not realistic in cost. Some barriers derive from industry structure, leading to contention among actors with adverse interests. And some barriers are grounded in human nature: malice, negligence, ignorance, or foolishness. We are concerned with a specific type of barrier that we deem a *harm*, which we define as *an impairment – either with respect to an individual, a firm or society – to an entity's welfare interests, relative to the normal expectations of the time and context.* We have taken this definition of harm from Kleining,[2] which carefully develops the concept of harm, starting with different historic traditions, ranging from an injury experienced by an individual (a too-narrow conception) to a legal definition of a harm, which is a proxy for an actual harm that facilitates government intervention. Kleining

1. David Clark and kc claffy, "An Inventory of Aspirations for the Internet's future" (2015), `http://www.caida.org/publications/papers/2015/inventory_aspirations_internets_future`.

2. Kleinig, "Crime and the Concept of Harm," *American Philosophical Quarterly* 15, no. 1 (1978).

reviews the relevance of concepts such as trust, and the relation of harms to the individual and harms to society. His view of a harm as an impairment relative to a welfare interest derives from a critical analysis of what those terms and competing terms such as well-being convey.[3]

We are similarly concerned with how to justify regulatory intervention to compensate for behaviors or activities that cause or threaten harm. Our definition of harm does not encompass any event that disadvantages an individual, business, entrepreneurial aspiration, or even entire industry. The natural selection aspect of market dynamics is a feature, not a bug. But, for example, when a company fails due to anticompetitive behavior of an incumbent rival, we consider the incumbent to have harmed the company and its stakeholders. Antitrust legislation and enforcement authorities exist to mitigate the harms of anticompetitive behavior. One motivation for this study is to elucidate other harms that have captured attention in today's Internet policy debates, and for which current regulatory apparatus seems ill-equipped to handle. This motivation will influence the structure of our taxonomy of harms.

More precisely, our goal is to outline the landscape of harms as a step to discussing the landscape of remedies. We are interested in situations where a policy-maker may find a harm of sufficient importance to develop and implement steps to remedy that harm. As such, we try to organize the list of harms in a way that sheds light on which actors have responsibility for allowing the harm, or should be tasked with preventing or mitigating the harm. For this reason, we prefer a structural classification of harms, rather than a consequential classification. We attempt to classify harms by layer or segment of the ecosystem in which they arise. And as with our conversation about aspirations, we want a framework that can help illuminate the interactions and tradeoffs – conflicting articulations of welfare interests – in attempting to mitigate any specific harm. We assume our list may not be complete. We also recognize that what constitutes a harm may change over time, such as what bandwidth is sufficient to consider a home broadband-connected. The continual evolution of the ecosystem is one reason that our definition specifies that a harm arises from an impairment *relative to the normal expectations of the time and context.*

We also want to consider how unique harms on the Internet are to the Internet ecosystem. Some harms are essentially similar to a corresponding harm that pre-dates the Internet; the only challenge may be defining and detecting its occurrence in a digital ecosystem, e.g., fraud or forgery. Other harms are well-known, but the accelerating and amplifying power of the Internet as a platform raises these harms to a level that triggers debate as to how regulatory intervention might mitigate the harm. Targeted advertising, surveillance, and risk of addiction to video games or television have been around much longer than the Internet, but when the scope, precision, and potency expand by orders of magnitude, we must acknowledge the potential for new harms. Still other harms are unique to a digital environment and the political economy in which it is embedded, such as attacks on networks and digital information.[4]

Governments and advocacy groups express many harms on our list as societal harms – undesirable outcomes for the citizenry, prevention and mitigation of which is in the public interest. And yet the Internet's architecture and infrastructure are now primarily under the stewardship of the private sector, driven by profitability and commercial viability, constrained by technological and economic circumstances, and sustained by interconnecting and interoperating competitors in a multistakeholder ecosystem. Navigating the inherent tension between private sector objectives and prevention of societal harms is essential to shaping the future of the Internet, as it would be for any sector where society has a strong public interest in a creation of the private sector.

One debate that this paper can trigger is whether a particular harm rises to the level that requires intervention. We hope the taxonomy can help structure these debates, which must balance classic tradeoffs between "freedom from" (e.g., harassment) vs. "freedom to" (e.g., free expression). A related question is how to measure the prevalence or extent of a harm, or the effectiveness of a remedy. An obvious and often repeated root cause of the broken state of public dialogue is the lack of data to inform debate. The path from harm to mitigation requires some approach to demonstrating the existence of a harm, and then debate

3. Kleinig, "Crime and the Concept of Harm" provides a number of citations to further critical analysis of harm, to which we refer the interested reader.

4. Anderson, Ross and Barton, Chris, and Bohme, Rainer, and Clayton, Richard, and van Eeten, Michel and Levi, Michael, and Moore, Tyler and Savage, Stefan, "Measuring the Cost of Cybercrime," in *Workshop on Economics and Information Science*, `https://www.econinfosec.org/archive/weis2012/papers/Anderson_WEIS2012.pdf` (2012).

about the merit of specific steps in preventing or mitigating the harm. Ascertaining whether any given intervention was effective requires not only longitudinal measurement, but also making a convincing claim that the intervention is what caused the effect.

Our network architecture background frames our thinking about harms in terms of layers of the Internet architecture,[5] and we attempted to begin conceptually at the lower layers, e.g., physical access, and work our way up. At the same time, the overarching concerns about the trustworthiness of information on the Internet led us to consider the classic *CIA triad* (confidentiality, integrity, availability) used in information security, and to classify harms according to their impedance to these three objectives.[6] These two framings are orthogonal, in that we can identify harms to confidentiality, integrity, and availability at different layers of the architecture. Importantly, harms at some layers of the architecture merit more regulatory attention, i.e., are a potentially larger threat to welfare interests than harms at other layers. We will start with this triad, in reverse order, because it allows us to cover a broad range of policy debates in order of their received attention thus far. It yields a rather coarse classification, but serves as one useful structuring tool for our analysis.

So, our first category will be harms to *availability* or access. This category includes policy goals such as *reach*, *ubiquity*, *affordability*, *uptake*, *unblocked access*, and *generality*, *evolution of performance* of the Internet, and the harms that result from not achieving them. We then shift our attention to a class of harms that impede our trust in the security of the Internet. From cybercrime, to unlawful behavior, to national security threats, the range of security-related harms on the Internet merits its own study, and will consume a significant fraction of this one. We organize this part of the discussion in two parts.

We first consider harms to the *integrity* of the user experience. In the information security domain, integrity generally refers to the assurance that data is not modified in an unauthorized or undetected manner. We can expand this definition in an Internet context to the idea that the user transmits and receives data they want to transmit/receive, and does not transmit/receive data they do not want to transmit/receive. This has some conceptual overlap with availability, but we think the delineation will be clear.

The third leg of the triad, *confidentiality*, brings our focus up the layers, from pipes to platforms, so to speak. Our discussion ~~but~~ takes us well beyond what information security experts mean by confidentiality to include the broader and less well-defined concept of privacy. We cannot describe harms to confidentiality of information and communications without covering the potential and actual harms arising from the emergence of powerful digital platforms, many of which use personal information about users as a primary essential input.

This acknowledgement of the age of "information capitalism" serves as a link to the second half of the paper, which focuses on more traditional economic, political, and social harms, and how they manifest in the Internet ecosystem. We begin with the policy statement issued by the U.S. FCC in 2005 adopting four principles: that consumers are entitled to choice in legal content, apps and services, devices, and broadband access service providers.[7] Their reasoning implies that a loss of choice, measured as insufficient competition, is a harm. However, the landscape of choice and competition – and the associated aspiration of innovation – is complex, and there are not easy criteria to judge when harm has occurred. Furthermore, several forces tend to militate against choice and competition in the Internet ecosystem. Market forces naturally push toward consolidation in any competitive industry (competition leads to winners), but global network effects,

5. kc claffy and David Clark, "Platform Models for Sustainable Internet Regulation," *Journal of Information Policy* 4 (2014): 463–488, ISSN: 21583897, `http://www.jstor.org/stable/10.5325/jinfopoli.4.2014.0463`.

6. The Federal Information Security Management Act (FISMA) defines the relation between information security and the CIA triad as follows: *(1) The term "information security" means protecting information and information systems from unauthorized access, use, disclosure, disruption, modification, or destruction in order to provide:*

A. Integrity, which means guarding against improper information modification or destruction, and includes ensuring information nonrepudiation, accuracy, and authenticity;

B. Confidentiality, which means preserving authorized restrictions on access and disclosure, including a means for protecting personal privacy and proprietary information; and

C. Availability, which means ensuring timely and reliable access to, and use of, information. (FISMA https://csrc.nist.gov/topics/laws-and-regulations/laws/fisma)

7. Federal Communications Commission, *FCC 05-151, Policy Statement*, In the matter of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, CC Docket No. 02-33, etc., `http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.doc`, 2005.

the marginal economics of digitization of commerce and culture, and even weaknesses in the underlying Internet architecture, amplify these forces on the Internet. We attempt to flesh out the nature of harms that derive from the consolidation and centralization of power and capital, concluding slowed macroeconomic growth, productivity, and innovation, political polarization, the death of trusted journalism, and corruption of democratic political processes.

In each section we include a discussion of what measurements would help to shed light on the extent of a harm, what general sorts of remedies might be effective to mitigate the harm, and what actors are in a position to undertake measurement and mitigation. Where possible, we will organize these considerations starting with the layer or segment at which a harm arises, which should illuminate which actors are responsible for the behaviors that cause the harm, or actions that might mitigate the harm. The actor that can best mitigate a harm need not (and often will not be) the actor that causes the harm. The tradition of tort law is that the party that should mitigate a harm is the one that is in the best position to do so, not necessarily the actors that contribute to the harm. Both good engineering design and regulation should follow this principle. Thus, for example, the packet forwarding layer of the Internet may allow malicious parties to observe the traffic flows; it is the end-node that can best mitigate this harm by encrypting the traffic. The Internet may lose packets; having the end-points retransmit lost packets is a much better solution than demanding that the forwarding infrastructure never lose a packet.

As a first principle, we conjecture that it will be most effective to mitigate a harm at or above the layer at which it arises, although this rule is not always true. Typically, lower layers of the ecosystem manifest more generality, and the generality raises a challenge to the construction of a useful remedy to a harm. Either the remedy is narrowly constructed, in which case the generality of the lower layer allows the misbehaving actor to exploit that generality to evade the remedy, or the remedy is broadly constructed, in which case it does collateral damage to other users of that layer. If one seeks a remedy at a layer higher than that at which the harm is occurring, the increasing specificity of the higher layer may imply that the remedy is applicable only in a specific context, so different contexts may require different remedies. Thus, for example, applications can to some extent compensate for failures of resilience at lower layers of the Internet and mitigate the risk of observation of traffic in the network, and the routing protocols of the Internet are designed to compensate for the failure of links and routers. We emphasize that in most cases the definitional and measurement issues are fundamental and daunting – and we hope this document can help structure conversation on directions that the research community, policy makers, and funding agencies can pursue to raise the rigor of Internet public policy discourse aimed at mitigating, avoiding, or remedying these and/or some future refined set of potential harms.

2 Harms to Availability of Internet Access

Several aspirations we previously catalogued related to the availability and quality of access to the Internet:[8]

1. *reach*, i.e., universal service, that every residence would be reached with suitable broadband service,
2. *ubiquity*, in the context of wireless coverage,
3. *evolution of performance*, to enable use of evolving services,
4. *affordability*, to facilitate uptake,
5. *uptake* by users, so more of the population could benefit from broadband services,
6. *unblocked access* to allow user choice in content,
7. *generality* to enable innovation at the edge.

8. Clark and claffy, "An Inventory of Aspirations for the Internet's future."

As broadband (Internet) pervades every aspect of society, lack of access constitutes a harm as we have defined it: an impairment to one's welfare interests, relative to the normal expectations of the time and context. The U.S. government's National Broadband Plan in 2010 validated the concept of exclusion from broadband service as a harm: "All Americans should have access to broadband service with sufficient capabilities, all should be able to afford broadband and all should have the opportunity to develop digital literacy skills to take advantage of broadband."[9] The allocation of universal service funds to broadband, and the expanding number of efforts, subsidized or supported entirely by the public sector, to bring broadband to unserved and underserved areas, all signal that lack of access to broadband has qualified as a harm.

Another area of contention around this harm is the *evolution of performance* metrics such as bandwidth (is lack of universal 4K TV capacity really a harm?),[10] and the most cost-effective methods to achieve penetration at a given service level. Inferior broadband performance is a quintessential example of a harm that changes definitionally over time and even across countries. Two decades ago, there was no consensus that lack of residential broadband access constituted a harm, and there was no technology available to support gigabit residential access. At some point, gigabit access technology became available but not at a cost that justified widespread deployment. Some imagine a future point when gigabit deployment is practical for most users, and normal for many, even assumed by popular application developers as the norm. At such a point in time, society might judge that the loss of welfare to those who still cannot obtain gigabit broadband access is a barrier that rises to the level of a harm, justifying some remedy to mitigate the residual barrier.

Thus, the definition of minimal acceptable service evolves with technology and the basket of applications the general public tends to use. With respect to residential broadband access, the FCC, the OECD and the ITU use *peak speed* (sometimes download and upload bandwidths) to compare offerings across countries.[11] In 2015 the U.S. FCC redefined broadband access as having a minimum peak speed of 25 Mb/s download and 4 Mb/s upload, compared to a previous threshold of 4 Mb/s download and 1 Mb/s upload.[12] This redefinition caused major changes in the determination of how many households had access to broadband, and the degree of competition in the service. In contrast, the ITU and its partners, concerned with the developing world, defined broadband in 2016 as a service with at least 256 Kb/s download, 1% of the FCC definition.[13]

Other available metrics are equally or more significant: high traffic volume applications such as streaming video make *usage caps* a concern, and real-time interactive (voice, videoconferencing) applications bring attention to metrics of delay and its variance (jitter). Today, many providers in the U.S. set usage caps of 1000 GB/month,[14] which would allow no more than 330 hours a month or about 11 hours a day of Netflix HD video for all devices in the household combined.[15] Eleven hours a day of HD video may seem adequate, unless there are multiple viewers in the home. However, new applications are emerging that will greatly increase the demand for capacity. HD video will more than double the required data rate, so for 4K video, 1000 GB/month will allow a household to watch about 5 hours a day. These numbers reflect the high end of current applications—today it does not seem that a user that cannot watch as much HD TV as they want has been harmed. But is there a lower monthly usage cap that so disadvantages users, relative to the normal

9. Federal Communications Commission, *The National Broadband Plan: Connecting America*, http://download.broadband.gov/plan/national-broadband-plan.pdf, 2010, p. xiii.

10. See William Lehr and Douglas Sicker, "Would You Like Your Internet With or Without Video?," *Journal of Law, Technology & Policy*, 2017,

11. Federal Communications Commission, *International Broadband Data Comparison Report*, GN Docket No. 17-199, 2019.

12. Federal Communications Commission, *2015 Broadband Progress Report And Notice Of Inquiry On Immediate Action To Accelerate Deployment*, `https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-10A1.pdf`, February 4, 2015.

13. "Partnership on Measuring ICT for Development", *Core List of ICT Indicators*, March 2016, `https://www.itu.int/en/ITU-D/Statistics/Documents/coreindicators/Core-List-of-Indicators_March2016.pdf`.

14. BroadbandNow, at https://broadbandnow.com/internet-providers-with-data-caps, reports on data caps from all U.S. ISPs. The numbers are highly variable, but 1,000 GB/month is typical for the larger ISPs.

15. Netflix estimates that downloading HD video can consume up to 3 GB/hour. See https://help.netflix.com/en/node/87. Also, "The idea that latency and packet loss can be as important as bandwidth is not new. But it is one that plays almost no role in contemporary policy debates. This is one of the gravest mistakes we are making today. It is akin to having a transportation policy that focuses on miles of highway constructed but pays no attention to whether those highways actually decrease commute times or accidents." in http://www.techpolicydaily.com/communications/five-faulty-premises-part-2-overstated-need-broadband-education-healthcare/

expectations of the time and context, that we could consider a user facing that usage cap to be impaired or harmed? And if so, what response would be justified? Perhaps the ISP with the low usage cap is in a region that is very costly to serve. Or perhaps the ISP faces no competition, and chooses not to invest in more capacity.[16]

Another dimension of harm relates to the aspiration of *evolution of performance*. If, as a part of specifying technology suitable to receive universal access subsidies, governments specify minimum performance requirements but do not require that the technology be able to evolve over time to higher levels of performance, the result might be a remedy that is suitable now, but not in the future, thus recreating the harm of inadequate availability at a later time.

This challenge of cost is complex and ever-present. *Affordability* can be a barrier to access, but it is not obvious what price should be the target to mitigate the contribution of price to the harm. Rural living has many consequences with respect to cost of living, with some things less costly and some more. The real issue may not be the price of rural broadband, but that rural areas often have many residents with lower income levels. The challenge of overcoming adoption by the poor merits a different intervention than overcoming the cost of rural deployment.

Another aspiration related to access is *uptake*. Even with broadband available, some may choose not to connect. As essential social services migrate to the Internet to increase the efficiency of delivering them, non-users are increasingly disadvantaged. Pew's surveys of American adults show a steady increase in use of the Internet, from 52% in 2000 to 90% in 2019.[17] Their data shows a flattening of uptake around 2016, when 86% of those surveyed used the Internet.

If the non-users abstain by choice, it is hard to claim that there is any harm. But perhaps non-users are constrained by barriers to uptake that are beyond their control. The FCC's survey taken for the National Broadband Plan in 2009 found about 22% of Americans (U.S.) did not use the Internet, and identified a variety of reasons, including cost (36%), lack of skills (22%), and perception of insufficient value (19%).[18] A similar survey by Pew in 2013 reported that 34% of non-users thought the Internet was not important, 19% mentioned issues related to cost, and 32% mentioned issues related to difficulty of use.[19] In that report, Pew reported that 27% of those over 65 did not go online, while usage of younger Americans is essentially ubiquitous (100% of those surveyed in the age group 18-29, and 97% of those between 30-49). If barriers to uptake felt by the elderly (for example) are leading to an impairment to their welfare interests, there is cause to believe they are being harmed.

The interplay of cost and uptake is capability-dependent. Many people willingly pay ~~over~~ over $100/month for smart phone service, but do not subscribe to residential service, suggesting that people value ubiquity more than performance. In a 2019 report,[20] 17% of users with a smart phone did not have broadband at home. According to that report, the number of homes with broadband has held around 70% since 2013. In 2019, 75% of urban homes, 79% of suburban homes, and 63% of rural homes had broadband.

Other dimensions of Internet availability: resilience and reliability The discussion to this point relates to availability in a physical sense—the harm that the network does not reach certain regions. The second aspect of availability (which we could include here or in §3) is that the network stops working. The Internet suffers minor outages all the time, mostly for operational reasons. Nations take down their regions of the Internet for one or another political reason. Packets sent to a legitimate destination may not arrive for

16. There is also the question, with implications both for technology and policy, as to whether the delivery of high-quality video is actually a suitable application for the Internet. The demands for capacity and the implications for industry structure are largely specific to the media entertainment sector, and perhaps this sector should be segregated from the internet. For a detailed analysis of this proposition, see Lehr and Sicker, "Would You Like Your Internet With or Without Video?" Under what circumstances should the inability to watch high-resolution video be considered a *harm*?

17. "Pew Research Center", *Internet/Broadband Fact Sheet*, `https://www.pewinternet.org/fact-sheet/internet-broadband/`, June 2019.

18. Federal Communications Commission, *Broadband Adoption and Use in America*, http://transition.fcc.gov/DiversityFAC/032410/consumer-survey-horrigan.pdf, 2010.

19. "Pew Research Center", *Who's Not Online and Why*, `https://www.pewinternet.org/2013/09/25/whos-not-online-and-why/`, September 2013.

20. Center", *Internet/Broadband Fact Sheet*.

many reasons, including unintended flaws or insecurities in the routing system (see §3) or intentional blocking of traffic based on port, address, or content. There have been apparent attacks on network communications infrastructure such as undersea cables.

However, we have never had a massive outage of the Internet. An outage of sufficient size would rise to the level where the harm relates to our aspiration of *national security*. But we need not imagine a nation-state attack on the Internet to appreciate that as services traditionally covered by legacy telephony regulations to ensure reliability and availability move to packet-based technology, regulators will have to develop new operational definitions that capture meaningful constraints on the feasibility of such services on the public Internet. The most serious outages to our current communications infrastructure have been caused by natural disasters such as hurricanes and floods. The consequences for the public are very serious.

With respect to attacks, the probability of a massive failure may be low, but the consequences of a successful attack would be high. We know it is not practical to engineer network infrastructure so that the probability of a successful attack is zero. The best approach is to design these systems with a high degree of resilience, so that they can continue to function while the consequences of an attack are reversed. The challenge for the U.S., as it considers whether its Internet infrastructure is sufficiently resilient, is how to engage the many independent (and competitive) Internet Service Providers that make up the nation's Internet. In the era of the AT&T regulated monopoly, where the company enjoyed a guaranteed profit, AT&T was happy to work with the government to make the phone system highly resilient. Parts of the AT&T backbone were engineered to survive a nuclear attack. AT&T provided centralized network management and planning functions. In contrast, there is no central planning for today's Internet. Thousands of firms make independent decisions about configuration, interconnection, redundancy, extra capacity, etc. This reality makes clear why the first major challenge with respect to resilience is getting data that would inform a study of how resilient our nations's Internet is.

2.1 Measurements and remedies: challenges and opportunities.

For each category of harm, we will include a section in which we discuss what measurement might shed light on the extent of the various harms we identify, what sort of remedies may serve to mitigate the harms, and what actors are in a position to undertake measurement and mitigation. For this category, we discuss options for measurement and remedy of four harms: lack of availability, low quality service, high price service, and inadequate resilience.

2.1.1 Measurement and remedies to harms to availability

Of all the categories of harms we will discuss, lack of physical availability (what we called *reach* and *ubiquity*) is the one that seems like it should be most amenable to empirical measurement, perhaps because we describe it in the simplest numerical measurements, e.g., served households, access bandwidth, or wireless signal strength. And yet the statistics we quoted in this section illustrate the complexity of mapping impairments in metrics of infrastructure availability to harms.

There seems to be consensus that the broadband coverage maps derived from the current data collection practices of the FCC,[21] based on Form 477 filings from network operators, overestimates connectivity, especially in rural communities.[22] Microsoft has recently reported inconsistencies between the FCC-reported data and Microsoft's own observations of broadband availability.[23] There are recent efforts by the FCC to improve the quality of the data that they gather and report.

21. Federal Communications Commission, *2019 Broadband Deployment Report*, https://www.fcc.gov/document/broadband-deployment-report-digital-divide-narrowing-substantially-0, 2019.

22. Shiva Stella, *Public Knowledge Warns Broadband Report Not Based on Quality Data*, `https://www.publicknowledge.org/press-release/public-knowledge-warns-broadband-report-not-based-on-quality-data`, 2019.

23. Paul Garnett, *Better broadband data can lend a voice to rural Americans*, `https://blogs.microsoft.com/on-the-issues/2018/08/16/better-broadband-data-can-lend-a-voice-to-rural-americans/`, August 2018; John Kahan, Chief Data Analytics Officer, *It's time for a new approach for mapping broadband data to better serve Americans*, `https://blogs.microsoft.com/on-the-issues/2019/04/08/its-time-for-a-new-approach-for-mapping-broadband-data-to-better-serve-americans/`, April 2019.

With respect to wireless coverage availability, given an agreed definition of wireless service capability, one can construct maps of coverage, based on data from providers, user surveys or field measurements. But seamlessness of access is also hard to quantify, e.g., performance when moving across access points. Measurement of resilience, reliability, and QoS are challenging to sustain at scale.

Mack *et al.* observed that the "never-ending quest for more data, more precision, reliability, validity and consistency over time" in mapping efforts is challenged by uncoordinated changes in reporting requirements, technology, and operational definitions. Longitudinal analysis across inconsistent patchworks of historical data requires creative methods to synthesize and interpret existing imperfect data.[24]

To remedy the harm of inadequate access, the transition from sanctioned monopoly to hoped-for competition in last-mile access has changed the range of viable approaches. In the days of the regulated monopoly, the regulator could give the incumbent the duty of carrier of last resort, and obligate it to build out unprofitable areas, with the costs folded in to the regulated rates. Today, the best approach seems to be direct public sector investment. For example, to achieve connectivity in rural areas that lack sufficient revenue to justify private investment in infrastructure deployment, governments have provided subsidy or tax incentives to build or maintain networks. In some cases the public sector has directly funded construction. The U.S. has a long history of federal support for rural deployment of critical infrastructure that the private sector would otherwise not deploy, most notably the 1936 Rural Electrification Act, which created the Rural Electrification Administration to issue loans and other help to rural organizations setting up their own power systems. In the United States, direct public investment in broadband has happened at multiple levels, from federal stimulus money[25] to municipal construction of residential broadband networks.

2.1.2 Measurement and remedies to harms from low quality service

Service with inadequate bandwidth (or other impairments) can be sampled from the edge of the network, so third-party measurement can provide some perspective. The FCC established its Measuring Broadband America program to perform such measurements, using a sample of about ten thousand homes. If the broadband service is provided using government funding as a part of a remedy to lack of availability, careful specification by the government of the service requirements that qualify for funding can help mitigate this harm.

2.1.3 Measurement and remedies to harms from high price service

Data on price of service can be collected, perhaps with considerable effort and incompletely due to the difficulty of determining how costs differ by serving area and customer. Since the data is in principle public, requiring the providers to report their pricing does not seem unreasonable. If price is seen as a major barrier to uptake, some sort of subsidy program could in instituted, or even some sort of price regulation.

2.1.4 Measurement and remedies to harms related to resilience

With respect to the harms that might arise due to lack of resilience in the infrastructure, the measurement challenges are substantial. Because of the decentralized character of the Internet's infrastructure, any attempt to define and measure resilience will be challenging. However, there are probably ways to assess, at least to some extent, how resilient the Internet would be to various failures. The Border Gateway Protocol (BGP), which implements the Internet's global routing protocol, allows destinations to announce more than one route by which they can be reached. In principle, all the BGP routing information is potentially public—the challenge is to gather enough of it to discover what alternate routes might exist. If enough of those routing messages could be captured, it might be possible to construct a routing simulator for the Internet, and perform various tests on the system, somewhat like a stress test regulators now require of financial

24. Mack, Elizabeth and Dutton, William H. and Rikard, RV and Yankelevich, Aleksandr, "Mapping and Measuring the Information Society: A Social Science Perspective on the Opportunities, Problems and Prospects of Broadband Internet Data in the United States," `https://ssrn.com/abstract=3333292`, *The Information Society* 35 (2019), doi:`10.1080/01972243.2019.1574526`.

25. See the Broadband Technologies Opportunities Program, http://www2.ntia.doc.gov

institutions. As part of understanding the potential of a harm to network availability, it might be possible to define a level of stress test that an ISP must pass in order to achieve some sort of rating, or to be used for critical sectors of the industry. However, this approach would require the identification (or creation) of an organization with the capacity and authority to undertake this data gathering and sharing of the analysis, which would require the backing of the governments.

A stress test of resilience would be a revealing exercise, but it would not capture the full degree of resilience in the Internet, because there is probably a high degree of hidden resilience in the ecosystem. After the 9/11 attack destroyed major switching centers in Manhattan, operators physically reconfigured interconnections, creating paths that were not there before. The cellular industry has mobile base stations (Cell on Wheels or COWs) that they physically deploy after disasters as part of network recovery. There is no way that measuring the current Internet can give a complete answer as to how resilient the network actually is, because it cannot detect the potential for physical reconfiguration.

Today, providers of traditional telephone service are required to report outages of a sufficient scope to the FCC. A similar program could (and probably should) be instituted for Internet service providers. Outages can be measured by third parties (using some clever techniques and perhaps access to proprietary information), but confirmation by the ISPs would be valuable.

To remedy harms associated with network failures, the Internet Service Providers, who provision and activate the physical infrastructure, must be the ultimate actors, since they control the deployment and management of the physical infrastructure. If the government is concerned about physical availability or resilience, it must work with these actors. However, designers of applications can also take steps to improve availability, so long as a sufficient physical path to the user exists. As providers of content (and other sorts of applications) exploit CDNs and install more caches closer to the user, successful operation of the application depends on a smaller and smaller region. The government could encourage designers of applications to stress test their applications, to make sure the designers know what network assets they actually depend on. The observation that application designers can to some extent compensate for failures at a lower layer is an example of our point that the actor that can mitigate a harm is not necessarily the actor that controls the layer at which the harm occurs. The design philosophy of the Internet was that the lower layers were not expected to be perfect, but just to deliver a reasonable service (so-called "best effort" delivery), and applications would include approaches to compensate for lower-layer failures. Designers of applications that aspire to a high degree of resilience (consider a "911 app") will need to exercise innovative thinking to ensure that the application is sufficiently resilient as failures at lower layers occur.

Finally, with respect to resilience at the application layer, the fundamental measurement challenge is that applications today depend on assets other than the public Internet. In addition to cloud computing and CDNs, applications are more and more using communication links and networks that are not a part of the Internet. Independent researchers are not able to get any visibility into these networks, and governments are not going to get any visibility unless they compel relevant data collection. Only with this data can governments begin to make any independent judgments about application resilience. If the government cares about overall resilience of the Internet ecosystem, it will have to extend its data-gathering and monitoring beyond the Internet itself to the other assets on which today's applications depend.

3 Harms to the Integrity of the Internet Experience

We are wary of a term as broad as *integrity* when referring to use of the Internet, but more common terms such as *security* and *trustworthy* seem even broader. The strictest (information security) definition of a secure system is one that meets its specification, even if under attack. But meeting its specification does not make a system *trustworthy* – this latter property has broader scope. Users mistrust today's Internet for many reasons. Often, the Internet proves untrustworthy because in its design, in pursuit of some other goal, actors have deliberately implemented risky or untrustworthy behaviors into the ecosystem, which malicious actors then exploit.

Although many have contemplated and designed more secure network architectures, there is always a tradeoff – restricting capabilities comes at a cost to innovation, generality, deployability, and usability. We

remind the reader that our baseline definition of integrity is that the user transmits and receives data they want to transmit/receive, and does not transmit/receive data they do not want to transmit/receive. The data is not modified in ways that is not intended by the communicants. Because harms to integrity can happen at so many layers of the architecture, we work our way up from the bottom. We note some of these service-related harms reduce to *availability* harms of the CIA triad; for purposes we will distinguish availability of infrastructure, as in access to infrastructure covered in the previous section, with availability of services running on the infrastructure.

3.1 Harms occurring at physical and network layers

The physical and network layers of the ecosystem have a comparatively simple task—move packets from a source to the intended destinations. We consider three harms to the integrity of the network service at these layers: traffic intended for one destination is delivered elsewhere; traffic travels via a path that allows a third party to observe or modify it in ways not requested by the communicating parties; and end-nodes receive unwelcome traffic.

1. The first harm, when **traffic intended for one destination is delivered to a different destination**, can arise due to behaviors by multiple actors. A network can forward packets based on erroneous routing information (manipulation of the Border Gateway Protocol, or BGP), or flaws in the DNS or the TLS certificate authority system can mislead a sender into using the wrong destination IP address. It can be hard (impossible) to distinguish this failure mode from traffic not being delivered at all. That is, it is easy for an end-node to detect that traffic is not being delivered; it is hard to tell what else may have happened to it. In the language of information security, this outcome is also either a loss of availability, or (to the extent to which the alternative destination can usefully examine the traffic), a loss of confidentiality.

2. The second, related harm is when **traffic travels via a path that allows a third party to observe or modify it in ways not requested by the communicating parties.** If the traffic is encrypted, the receiving endpoint can detect tampering. Mere observation is much harder, sometimes impossible, to detect. The most sophisticated forms of this harm include either covert or overt manipulation of encryption keys or certificates to allow observation or modification of encrypted traffic. As an example, in July 2019, the nation of Kazakhstan mandated that all of its citizens install on their computer a government-issued *root certificate*, after which the government can intercept, inspect, block or modify (essentially) all encrypted traffic from those citizens.[26] This example requires active participation of the nation's ISPs to implement the interception. This approach sacrifices confidentiality, integrity, and availability of information in the country to a higher cause: national security as the government defines it. It is an excellent illustration of how words such as "security" or "integrity" are a function of values that differ across governments.

3. The third and final harm in this category is that **end-nodes receive unwelcome traffic.** Examples range from port scans to massive DDoS attacks. The lack of network layer source address authentication in the IP architecture exacerbates this harm: for many classes of this harm, neither the target nor the ISP serving the target can tell where the traffic is coming from.

Indeed, all three fundamental functions of the Internet architecture: naming, routing, and addressing, all occur by default with no authentication on the validity of the source of information. A global network architecture that supported identity-tracking would require some authority (or set of authorities) to issue and validate those identities. Anonymity would also be impossible in such an architecture. Although anonymity at the network layer results in harms to society, there are perhaps greater risks of harms in

26. This situation is fluid and may change before the paper is published. See https://bugzilla.mozilla.org/show_bug.cgi?id=1567114 for a discussion of the technical and political aspects of the mandate, and https://www.zdnet.com/article/kazakhstan-government-is-now-intercepting-all-https-traffic/ for a recent report.

removing anonymity altogether.[27] Identity has its proper place on the Internet, when parties need to know their respective identities. In the TCP/IP architecture – and it is a common view in network architecture research – managing identity is an application-layer function, not a function of the lower layers. Application designers thus must carefully consider issues related to identity, which may not be their highest priority as they bring a new service to market.

End-nodes can to some extent mitigate these harms. Encryption is effective unless the adversary has the skills or power to corrupt the system. One can imagine additional mechanisms that allow an end-node to confirm certain aspects of identity—tools to cross-validate the DNS or the Certificate Authority system. However, such checks often lead to termination of the communication, to prevent the harm of observation or modification. This outcome is a failure of availability, which may equally suit the needs of the adversary. An ISP may be able to identify (and drop) some classes of malicious traffic,[28], but in general, the network cannot ascertain whether traffic is harmful – this decision must occur at a higher level.

3.2 Harms occurring in edge devices and routers

We enumerate five harms typically associated with end user devices or routers attached to the Internet: (1) a node hardware or software fails or is otherwise unstable; (2) a node allows unauthorized access from the network; (3) a node blocks installation or execution of non-malicious software authorized by the user; (4) the software update system downloads a malicious version of the operating system; or (5) the operating system, by design, imposes a harm on the user, e.g., tracks and reports user location to unauthorized third parties.

1. **An edge node hardware or software fails or is otherwise unstable.** This harm does not normally result in regulation, just public criticism. But in the case of sector-specific critical systems (health care or avionics) there are regulations that address system reliability.

2. **An edge node allows unauthorized access from the network.** A great deal has been written about edge nodes that are vulnerable to penetration via the network. We do not attempt to construct a detailed taxonomy of those vulnerabilities here, which can range from simple attacks that allow a penetrator to issue commands on the computer to more sophisticated attacks that install malware on the computer. It is not clear that edge-node software can always block unauthorized installation of software, nor can an edge node detect in general that a particular item of software is malicious. Today, there is no requirement that an edge node take steps to block installation of unauthorized software. However, in the future the state of system security may reach a point where systems that do not attempt to block the unauthorized installation of software may be given a lower score in a quality rating system, or even banned in certain circumstances.

   This harm takes a particularly pernicious form when the downloaded software turns the infected computer into part of a botnet, an often-massive collection of machines under the control of a so-called botmaster. Botnets are used for a variety of malicious purposes, including spam campaigns and distributed denial of service (DDoS) attacks.

3. **The system prevents the installation of desired applications, or an installed application will not run.** The provider of the system, perhaps via an app store, has blocked the app, or because the app depends on features not present in the system e.g., a video camera. These harms are probably mostly at the nuisance level for the user, but may be existential for the provider of the app. So the basis for determining whether a harm merits mitigation depends on whether the policy goal is to protect the competitors, protect the overall level of competition, or protect the consumer from consequential harm. See our discussion of innovation in §5.

4. **The software update system is corrupted such that a user downloads a malicious version of the system.** Providers of operating systems and other software packages take great care to protect

27. David Clark and Susan Landau, "Untangling Attribution," *Harvard National Security Journal* 2 (2 2011).
28. See, for example, Philipp Richter and Arthur Burger, "Scanning the Scanners: Sensing the Internet from a Massive Distributed Network Telescope," to appear (Internet Measurement Conference (IMC) '19, October 2019).

against this harm, because if a provider of software loses control of the update process the result is catastrophic. The legitimate provider of the software can lose control of the update system to an adversary, and may not be able to regain that control.

5. **The operating system of the end-node, by design, imposes a harm on the user.** This open-ended harm resembles harms at the IP layer: the application fails to run, or transfers (or lets apps transfer) data to third parties without permission, etc. An operating system of a mobile device could track the location of the device (and its user) and report that location to unauthorized third parties.

This last example illustrates the distinction between a taxonomy based on structural considerations as opposed to a taxonomy based on consequences. One could list unauthorized release of personal information (PII) as a harm, independent of how it happened. Laws could make this harm illegal. But to prevent the harm, one must understand in what part of the system the harm can arise–which actors allow the release, and which actors exploit the release. In the case of release of PII, the harm can arise in several parts of the system, each of which implies a different focus for prevention or mitigation.

An extreme form of system penetration, which rises to the level of a harm to the aspiration of *national security*, is that the Internet might be used as an attack vector to disable some other component of the nations's critical infrastructure: the power grid, the water supply, the supply system for oil and gas, and so on. This situation is complex, with many considerations contributing to the final risk profile. However, a simple remedy that at least raises the effort of the attacker is to disconnect these critical systems from the Internet. Today, there are many options that can implement the communications needs of critical infrastructure providers. It is no longer necessary to use the Internet itself.

Another consequence of a system penetration is that it may lead to theft of information–a data breach. Data breach seems like such an important harm that it is worth listing it separately. However, there has been much work on this problem, including approaches to mitigation (See §3.5.2).

Another potential risk, this to *national security*, is that edge devices and routers could be vulnerable to a targeted attack on infrastructure. If a hostile actor could render many routers on the Internet somehow inoperable, so that some part (e.g, the memory used to boot the router) needed to be physically replaced, the damage might persist for weeks. A related risk is the fact that a few vendors supply the equipment that makes up the majority of the commercial Internet, leaving a *monoculture risk*, where an attacker finds a vulnerability that can affect many machines at once. It is not clear how realistic this risk is: even if machines such as routers come from a small number of vendors, there are different hardware variants, and many releases of software. However, there are also benefits to monoculture:[29] reducing the complexity of system management, which allows more attention to details of management that can reduce the probability of a successful attack. (Put all your eggs in one basket; watch that basket..) As with all security strategies, it is a balance of risks vs. costs.

3.3 Harms occurring at the application layer

The range of harms at this level may seem unbounded. Our goal is to find baskets of harms that are complete but that suggest common modes of mitigation. If we cannot aggregate and consider classes of harms, we will be doomed to an endless game of whack-a-mole against every new variant of malicious behavior that creative people invent. We divide harms at the application layer into three broad categories: the application designer is malicious; the application designer has interests that are adverse to the user's; and the design of the application opens up the user to unregulated interaction with malicious parties. We examine each of these in turn.

1. We define a **malicious application** as one that implements behaviors that are a) undocumented, b) not authorized by the user, and c) yield harmful outcomes. The first two criteria for malware – undocumented and unauthorized behavior – should be reasonably clear. The last category may be

29. Daniel E. Geer, "The Evolution of Security," *Queue* (New York, NY, USA) 5, no. 3 (April 2007): 30–35, ISSN: 1542-7730, doi:10.1145/1242489.1242500, http://doi.acm.org/10.1145/1242489.1242500.

subjective in the limit, but in many cases is unambiguous. Attempts to prevent or mitigate malware thus far have fallen short of society's expectations. If the user is misled into installing the software, then the machine has not literally been accessed without permission, but there is clearly risk (or expectation) of a harm.

One form of malware is a "key-logger", which can capture passwords as they are typed. This harm can lead to identity theft, and various form of fraud.

Another highly harmful example of malware is ransomware, where the malware encrypts the contents of the attacked computer, making it unavailable until the victim pays a ransom. There are a number of mitigations that have been proposed, including specialized monitoring software installed on an end node that makes a backup copy of any file that is modified, and looks for patterns of modification that look like a malware attack. Once the attack is halted, the software can roll back the system to its pre-attack state.

Prevention of the harms of malware by technical means has proven difficult. Operating systems may include tools, e.g., sandboxing, to mitigate potential harms of installed malware. Clever attackers seem to find ways around protections, so the operating system cannot reasonably assume the total burden for preventing this harm. Security review and curation by app stores can limit these harms to some extent. However, harms from malware persist. Moreover, a user need not even download malicious code to suffer harm. If a user interacts with a malicious application over the network, the potential for mischief may be more limited, but it is also harder to detect what the application is doing, and to judge whether harms are occurring. The result is a situation where users, if they are aware of the problems, live in uncertain fear about possible harms, which erodes trust in use of the Internet at all.

2. **Operators of dominant applications add unwelcome behaviors, terms or conditions to the application.** Such behaviors (if documented) would not qualify as malicious, but may still cause harm. If an application is dominant, and users are feel a strong need to persist with the use of the application, the potential for this harm will increase. Examples would include rent-seeking if the application is fee-based, or increasingly intrusive tracking of users and precision targeted advertising.

   This last harm arguably fall under (in the U.S.) the FTC's enforcement of truth in advertising: an application should disclose all material consequences of executing it. Mandated transparency is an important building block in understanding what an application is doing, whether it conforms to its disclosures, and the potential for an actual harm. However, the goal of transparency (as an application is actually being used) is in tension with confidentiality. Should users (or technically informed advocates for users) have the ability to see and understand what their applications (or IoT devices with embedded applications) are sending over the Internet, in order to identify potentially malicious activity? Or should the app designer have the right to prevent such observation by encrypting the traffic from the application, justified in order to prevent unauthorized disclosure of sensitive user data? We expect (and argue for) acceptance of the appropriateness of mechanisms that allow the encryption to be broken under controlled circumstances—for example when the device is under the physical control of its owner/user. A typical user would not be expected to study and understand what application elements are transmitting, but researchers and product evaluators need data that provides a starting point for their research. Recent advocacy[30] and research efforts are tackling this challenge.[31]

3. **An application, by design (that is, not due to faulty implementation) includes risky modes of operation or interaction, without taking steps to mitigate the risks of these behaviors.** Applications that do not support interaction among multiple users do not allow interaction with malicious parties. Unless a flashlight app is itself malicious as defined above, it is not likely to

30. Keith Winstein, `https://www.politico.com/agenda/story/2015/06/internet-of-things-privacy-concerns-000107`, *Politico*, 2015,

31. Wilson, Judson and Wahby, Riad S. and Corrigan-Gibbs, Henry and Boneh, Dan and Levis, Philip and Winstein, Keith, "Trust but Verify: Auditing the Secure Internet of Things," in *Proceedings of the 15th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '17 (2017), `http://doi.acm.org/10.1145/3081333.3081342`.

trigger material harms. Applications that allow controlled interactions among known parties might trigger harms in principle, but are not likely to have serious risky modes. The relative lack of risk is due both to their limited functionality (e.g., no downloading of external software) and limited set of communication endpoints (known parties).

The class of applications most likely to permit material harms are those that embody the original vision of the Internet—applications that allow relatively unconstrained interaction among parties that do not know each other. That vision, combined with the rich affordances of many application, provide both powerful freedoms and risks of serious harms, e.g., spam, fake news, trolling or fake reviews. Applications with this character have some (usually implicit) social contract among the participants, but that social contract, in many cases, is neither clearly stated nor well enforced.

It is not reasonable to expect an application designer to identify in advance all patterns of risky use. However, providers that do not take steps to mitigate harms once they are occurring, and thus allow continuing harm to the overall user community, may in the future be the target of regulation.

3.4 Other taxonomies of harms to integrity

We know of two other recent efforts to catalog the harms to the integrity (security and trustworthiness) of the Internet experience. One ambitious effort yielded a U.K. government-commissioned report that called for the creation of a new regulatory agency to deal with them.[32] The scope of the harms in this report is: "...companies that allow users to share or discover user-generated content or interact with each other online." In other words, this list of harms fits into the third of our three broad application-layer categories. Their list of harms includes:

- Child sexual exploitation and abuse.
- Terrorist propaganda and recruitment.
- Glamorizing gang life.
- Content illegally uploaded from prisons.
- Sale of opioids and other illegal drugs.
- Anonymous abuse.
- Cyberbullying.
- Facilitating self-harm and suicide.
- Underage sharing of sexual imagery.
- Online disinformation.
- Online manipulation.
- Online abuse of public figures.

This list illustrates the wide range of harms to individuals that might fit into this overall class, and can frame debate about the exact form and justification for each example. Different nations might generate a different list, and prioritize concerns differently.

Another British research study – from the University of Oxford – taxonomized harms from the perspective of organizations,[33] although it also includes classes of harms to individuals who make up the organization. The focus was not on how a regulator might prevent or mitigate harms, but on how organizations should prepare to deal with the consequences of cyber-attack. Thus, their taxonomy looks quite different. Figure 1 reproduces the full diagram of harms they include in their taxonomy.

- Physical or digital harm.
- Economic harm.
- Psychological harm.
- Reputational harm
- Social and Societal harm.

A list of this sort might inform regulatory process, e.g., risk assessment for cyber-insurance. But knowing that a harm is economic is not helpful if the goal was to prevent the harm in the first place. Our taxonomy takes a more structure (architectural) approach, in hopes of more clearly illuminating the remedy landscape.

3.5 Measurements and remedies: challenges and opportunities.

We have been struggling with definitional and measurement challenges of Internet security since the first software we would now consider malicious was widely released on the Internet in 1988.[34] We take each subcategory of harms in turn.

3.5.1 Measurement and remedies to harms occurring at physical and network layers

As a reminder, we considered three harms at these layers: traffic intended for one destination is delivered elsewhere; traffic travels via a path that allows a third party to observe or modify it in ways not requested by the communicating parties; and end-nodes receive unwelcome traffic.

With respect to the harms of **mis-delivered or mis-routed traffic**, measurement to assess this landscape is challenging but possible. We discussed in §2 the possibility of an organization undertaking a systematic study of the Border Gateway Protocol (BGP) to explore issues of resilience. That effort could also detect abuses in that system, although that exercise would require several data sources, and cooperation of many ISPs in the Internet. As Internet infrastructure researchers for decades, we do not underestimate the challenge here: a highly-skilled organization would need to undertake this activity, which would require the backing of the government, or governments, depending on scope of the study.

We also described the DNS and TLS certificate systems as potential vectors of attacks on the network layer. A suitably skilled and backed organization could similarly undertake a measurement study of such abuses of the DNS. As with BGP, the academic research community performs studies of the DNS, but their work is hampered by the scale of the systems and lack of access to relevant information. In the case of the DNS the data access issues are even more diffuse and challenging than with the routing system. The proposition of compelling relevant DNS information would have to content with the failure of the multistakeholder ecosystem to make sufficient DNS-related data and resources available to yield scientific study of the DNS that can support effective remediation of harms and vulnerabilities.

The TLS Certificate Authority system is an interesting example of how different actors have moved to effect technical remedies in the face of a diffuse and seemingly intractable coordination problem. The CA system is highly distributed, with many independent actors, some malicious or with adverse interests. This situation would seem to call for a highly distributed coordinated response. However, Google, a centralized and highly powerful actor, devised a scheme to mitigate some of the harms that were occurring in the CA space. The scheme, called Certificate Transparency, uses a set of public registers of certificates, so that certificates issued by different authorities (whether legitimate or perhaps malicious) must be available for public inspection. Because they control the popular Chrome browser, Google was able to unilaterally impose this scheme on the Internet community. This exercise of power can trigger resentment in the community,

34. The Morris Worm, see https://en.wikipedia.org/wiki/Morris_worm.

Figure 1. Taxonomy of organizational cyber-harms.

Figure 1: *Taxonomy of Harms, taken from Figure 1 of Ioannis Agrafiotis et al., "A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate,"* Journal of Cybersecurity*, 2018,*

but also illustrates how a single actor (with both power and skills to innovate a solution) can transform the ecosystem.

End-nodes can defend themselves from the harm of misrouted traffic, but only to some extent. Encryption prevents an attacker who has captured the traffic from making use of it, except in the case of highly skilled attackers who subvert the encryption system. An end-node can tell if traffic is not being delivered, but it is not clear, given the design of the Internet, what such an end-node could do about it. One of us (with co-authors) proposed that the Internet needed to be augmented with a system we called the *knowledge plane*, which allowed nodes at the edge of the network to query why their actions are failing, and to send complaints to the correct actor.[35]

3.5.2 Measurement and remedies to harms occurring in edge devices and routers

Our five harms related to end user devices and routers were: (1) a node fails or is otherwise unstable; (2) a node allows unauthorized access from the network; (3) a node blocks installation or execution of non-malicious software authorized by the user; (4) the software update system mistakenly downloads a malicious version of the operating system; or (5) the operating system imposes a harm on the user, e.g., tracks and reports user location to unauthorized third parties.

With respect to the second harm – *system penetration* – the problem has long been understood, and the security research community has invested a great deal of effort cataloging possible approaches to preventing the harm, which we do not attempt to recapitulate here. The law contemplated this harm early on. The 1984 Computer Fraud and Abuse Act (CFAA)[36] lists as one of its illegal activities: *18 U.S. Code §?1030(a)(5)(A) knowingly caus[ing] the transmission of a program, information, code, or command, and as a result of such conduct, intentionally caus[ing] damage without authorization, to a protected computer*, where *damage* is defined as *any impairment to the integrity or availability of data, a program, a system, or information*. However, the complexity of investigation and prosecution of this crime, especially across jurisdictional borders, has been an obstacle to the effective use of this law as a remedy.

The CFAA also makes use of ransomware per se an illegal act: *(a)(7)(C) demand[ing] or request[ing] for money or other thing of value in relation to damage to a protected computer, where such damage was caused to facilitate the extortion*. Some states have passed laws making possession of ransomware illegal, but these laws raise definitional issues, because anyone that is successfully attacked has a copy of the software on their computer.[37]

We identified two specific harms that could arise as a consequence of a system penetration: recruitment of a machine into a botnet, and data exfiltration. With respect to **botnets**, there has been a lot of pragmatic experimentation as to the best approach to mitigate the harm. Technical intervention seems difficult, because the generality of the Internet gives the botmaster a broad spectrum of options to evade interdiction. Legal prosecution has been somewhat effective, but the cross-border nature of most botnets makes legal proceedings complex, time-consuming and costly. The community of people who are concerned with mitigating botnets are still looking for an effective counter-measure.

As mentioned in §3.2, one consequence of a system penetration is that it may lead to **data breach**. But data breach is a complex event, potentially involving many actors, and we believe it merits its own category of harm. There has been substantial work attempting to quantify the prevalence and cost, and on the best approaches to mitigation of this harm. Looking at credit card fraud as a specific objective of those who steal data, there are many actors involved in the credit card industry, and thus many actors that might bear responsibility to mitigate data breach. Converting credit cards to chip and pin, which made it harder to exploit stolen credit card data, was probably more effective in limiting harm from that form of stolen data than any degree of system hardening.

This example illustrates our point that it is often helpful to examine a complex harm from several perspectives, including asking what the real harm is—in this case system penetration, the actual theft,

35. David D. Clark et al., "A knowledge plane for the internet," in *In Proceedings of ACM SIGCOMM* (2003).

36. https://www.law.cornell.edu/uscode/text/18/1030

37. A news article discussing the illegality of ransomware can be found at Alan Neuhauser, "Can the Law Stop Ransomware?," `https://www.usnews.com/news/national-news/articles/2018-04-13/can-the-law-stop-ransomware`, 2018,

or the use of the data to commit fraud. Also, our understanding of the prevalence of data breach was greatly facilitated by various state laws that required that holders of personal data disclose incidents of data breach—an excellent example where mandating disclosure is the only effective way to measure a class of harm.

With respect to other harms we associated with the end-node, today there is little regulation of either the quality or the functionality of end-nodes. For consumer-grade software, the future will continue to be *caveat emptor*, but governments may move to impose regulation if harms become serious, such as IoT devices causing harm to physical property or life. Some harms, such as release of PII by an operating system, could be outlawed by privacy legislation, or disciplined by the FTC as an unfair practice.

3.5.3 Remedies to harms occurring at the application layer

Our basket of application-layer harms, which we acknowledged was effectively unbounded, included those that derive from: malicious applications; applications that allow gathered personal data to be stolen or otherwise unacceptably used or released; operators of dominant applications add unwelcome behaviors, terms or conditions to the application; and applications that by design include risky modes of operation or interaction, without taking steps to mitigate the risks of these behaviors.

With respect to **malicious applications**, the primary challenge is to detect the malicious behavior. If such behavior can be confirmed and is "obviously" malicious, a combination of law and simple banning can probably control the proliferation. However, it requires a high level of skill in some cases to detect what an application is actually doing, especially (as we discussed above) if all the communication between an application on an end-node and other points on the Internet are encrypted. The researcher studying such an application is in the position of an intelligence agency trying to sort out what an adversary is doing looking only at encrypted information. These sorts of efforts are costly, require a high level of skill, might in some cases be deemed illegal, and must be repeated for each application under consideration. A foundational question in this area is to discover where there are more methodical ways to understand the scope and manifestations of this harm.

With respect to **applications that impose abusive conditions** or otherwise have adverse interests, we return to this issue in §4. However, there is a gray area between malice, adverse interest, simple mis-design, and acceptable intention that will attract attention, and behavior in this gray area will be hard to categorize, measure and discipline. Consider a search engine that returns results about political candidates that seems to bias voters. Is such search engine behavior malicious, or acceptable political speech? Should it be detectable as an act of political speech? One study suggests that such search engine bias is not only effective but hard to detect.[38]

The category of applications that permit **malicious interactions among participants** is one of the most important harms to integrity of the Internet experience. It is also the most complex, with many specific examples, and raises serious challenges with respect both to measurement and remedy. The actors that will have to implement the remedies, however defined, are the application designers themselves. These designers will have to be nudged to action, either by public outcry or the imposition of government regulation, which will be complex to craft, especially with respect to activities that trigger First Amendment concerns.

In understanding how harms at the application layer may arise, and the possibility of mitigating them, a useful distinction is whether the application is centralized (under the control of a single operator) or decentralized. To the extent an application is centralized, it is reasonable to burden the operator with responsibility for mitigating harms that occur. But with a decentralized application, such as email or the web, there is no single operator. Either the design of the application – or the platform on which the decentralized nodes operate, for example, the Web protocols – must block malicious behavior by limiting the functionality of the application, which may not be acceptable from a utility perspective, or some third party must rate the decentralized elements participating in the service, based on whether they carry out

38. Robert Epstein and Ronald E. Robertson, "The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections," *Proceedings of the National Academy of Sciences* 112, no. 33 (2015): E4512–E4521, ISSN: 0027-8424, doi:10.1073/pnas.1419828112, eprint: https://www.pnas.org/content/112/33/E4512.full.pdf, https://www.pnas.org/content/112/33/E4512.

malicious behaviors that might cause harm. For example, a rating service might give a poor score to an email server that sends spam, or a search engine might rate web sites and return a warning as part of the research results saying that a particular web site is untrustworthy. As well, in a highly distributed system, every participant in the system (both human and application) must take action to protect themselves from malicious interactions. The need to cope with malicious actors within a distributed system puts a high burden on the designers of the protocols that bring a distributed application into existence, and is one of the forces moving the Internet toward centralized designs. (In addition to the lack of any legal accountability for harms, and consolidation forces that remove options for consumers.)

3.5.4 Remedies to harms: scientific research to support

We must acknowledge barriers to even studying harms in the Internet ecosystem, due to its largely proprietary commercial constituent components. We offer two examples where we anticipate regulation will be required to improve data integrity to support research.

The first area is macroscopic security metrics. Although thousands of Internet security companies exist today, there are no metrics by which we could claim that security is improving on the Internet. As an example, the threat intelligence community generates hundreds of sources of data on threats, but we have little understanding of the extent to which this data supports improvements in operational security. Available evidence is cause for skepticism that existing threat intelligence data is achieving its purported goals.[39] There is little correlation between different feed properties, so little ability to meaningfully differentiate data sources, and to prioritize certain metrics based on their specific need.

The second example is measuring the costs of cybercrime. In 2012, Anderson *et al.* undertook an effort to do so.[40] They distinguished three classes of crime: traditional crimes now conducted online, e.g., fraud, crimes amplified substantially by the Internet platform; and crimes enabled by the Internet, which they call *platform crimes*, such as the provision of botnets which facilitate other crimes rather than being used to extract money from victims directly. They observed that indirect costs of cybercrime include loss of trust in online transactions, missed business opportunities, and reduced uptake of online services. While not amenable to measurement, these generalized harms are a drag on society, and in particular on the online experience. They observe that for more traditional property crimes against individuals (such as car theft) the balance between the direct costs of the crimes and the indirect costs of deterrence is somewhat stable; this sector of crime and deterrence is "mature". In contrast, in online crime, where the overall landscape of crime is not yet well-understood, and the character of the crime evolves rapidly, there is not yet a well-understood approach to deterrence. Based on their estimates of costs, and the effectiveness of different approaches to deterrence, they argued that society would be better off spending less on technical defenses and more on catching crooks and putting them in jail.

We agree that control of unlawful activities is always a balance between the need to identify and punish miscreants, while avoiding excessive intrusion into the lives and activities of the law-abiding users. An over-zealous swing toward strict monitoring, weakened encryption, and identity tracking could represent a harm in itself. Perhaps clever technical design can somehow improve our ability to pursue miscreants, without generating harms to the public.[41] But the primary challenge of this recommendation is rooted in political economy: the first approach – technical defenses – enjoys the same scaling amplification offered by the software-based crime itself, and thus the same opportunities for return on investments to capital – hence the thousands of cybersecurity companies selling intrusion detection software-based products covered by intellectual property protection. The second approach has effectively a negative financial return on investment – it is a public sector service activity funded by taxpayers, and this public sector currently cannot compete with private sector salaries from the same skills. Indeed, the global nature of these crimes

39. Vector Guo Li et al., "Reading the Tea leaves: A Comparative Analysis of Threat Intelligence," in *28th USENIX Security Symposium (USENIX Security 19)* (USENIX Association, 2019), `https://www.usenix.org/conference/usenixsecurity19/presentation/li`.

40. Anderson, Ross and Barton, Chris, and Bohme, Rainer, and Clayton, Richard, and van Eeten, Michel and Levi, Michael, and Moore, Tyler and Savage, Stefan, "Measuring the Cost of Cybercrime."

41. Stefan Savage, "Lawful Device Access without Mass Surveillance Risk: A Technical Design Discussion," *ACM Computer and Communications Security (CCS)*, 2018,

makes investigation and prosecution even more expensive than for local crimes such as car theft. Many investigations cross jurisdictional boundaries, and technology cannot directly solve the cross-jurisdictional challenges. If the legal issues associated with cross-jurisdictional investigations persist, we may see the Internet drift toward Balkanization, where cross-border flows are more heavily regulated, inspected, or limited than flows that stay within a single jurisdiction. The EU General Data Protection Regulation has already created tension in this direction.

4 Harms to Confidentiality and Privacy

The information security meaning of *confidentiality* is stricter than what we will cover in this category – it refers to preserving authorization restrictions on access and disclosure of data. In everyday English, confidentiality tends to refer to an ethical duty to keep private any personal information shared with an attorney or doctor, unless the client explicitly consents to disclosure. We covered this violation of expectation in §3. *Privacy* is a broader term, but more often the term used in discussions of the harms of inappropriate data disclosure, or even data gathering and analysis, in the Internet ecosystem. Despite a wealth of scholarship on privacy, we have found a general failure to evaluate harms to privacy in a rigorous manner.[42] But privacy is not the thing harmed. the individual or organization is (perhaps) harmed. How?

We are not alone in finding a lack of clarity in this space. Alan Westin, in his 1967 book *Privacy and Freedom*, begins with the following provocative sentence: "Few values so fundamental to society as privacy have been left so undefined in social theory or have been the subject of such vague and confused writing by social scientists."[43] The logic that links basic concepts of privacy to the articulation of specific harms is weak. We conclude that the role of privacy as a concept is not to derive the harms, but to provide a legal and regulatory framework to prevent and/or remedy these harms.

Westin offers his own definition, rooted in self-determination and autonomy: "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others." This formulation is effectively the same that Warren and Brandeis posed in their landmark 1890 article, when they described the harm induced from loss of privacy as the making public facts about an individual's private life that the individual chooses to keep private.[44]

In that (now 130-year old) article, they stated as a principle that: "The right to privacy ceases upon the publication of the facts by the individual, or with his consent."[45] The word "publication" implies an action by the data subject to disclose the information, as opposed to the actions of a data collector observing the behavior and actions of the subject. The world is probably overdue for an update to this article, as the online world is such a different place from the world they wrote about. Users publish facts about themselves on Facebook, which raises the question as to whether a user should have rights of privacy with respect to what they post, if even intended only for friends.

Warren and Brandeis' reference to consent has led to the current U.S. focus on notice-and-consent as a basis for determining whether a harm has occurred. That is, the question of harms is legally resolved by the terms and conditions of the use of data-driven services. Bellovin recently reviewed legal developments building on notice-and-consent,[46] noting that even in the 1960s, legal scholars warned about relying too much on consent as a way to protect privacy. Bellovin quotes legal scholar Arthur R. Miller's 1967 testimony before a Senate subcommittee: *" Excessive reliance should not be placed on what too often is viewed as a universal solvent: the concept of consent. How much attention is the average citizen going to pay to a governmental form requesting consent to record or transmit information? It is extremely unlikely that the full ramifications of the consent will be spelled out in the form; if they were, the document probably would be so complex that the average citizen would find it incomprehensible. Moreover, in many cases the consent will*

42. We are not privacy scholars, any more than we are legal scholars. Nor do we attempt to produce here a complete literature review on privacy.

43. Alan Westin, *Privacy and Freedom* (IgPublishing, 1967).

44. Samuel Warren and Louis Brandeis, "The Right to Privacy," *Harvard Law Review* 4 (1890).

45. Ibid., pg 218.

46. Steven M Bellovin, *Comments on Privacy*, November 2018, doi:10.31228/osf.io/5s2vt, osf.io/preprints/lawarxiv/5s2vt.

be coerced, not necessarily by threatening a heavy fine or imprisonment, but more subtly by requiring consent as a prerequisite to application for a federal job, contract, or subsidy."[47]

In 1973 a U.S. advisory committee issued a set of principles that became known as the Fair Information Practice Principles (FIPP):

1. *There must be no personal data record keeping systems whose very existence is secret.*
2. *There must be a way for an individual to find out what information about him is in a record and how it is used.*
3. *There must be a way for an individual to prevent information about him that was obtained for one purpose from being used or made available for other purposes without his consent.*
4. *There must be a way for an individual to correct or amend a record of identifiable information about him.*
5. *Any organization creating, maintaining, using, or disseminating records of identifiable personal data must assure the reliability of the data for their intended use and must take precautions to prevent misuse of the data.*

The following year the U.S. government passed the Privacy Act of 1974, which implemented these principles for the federal government. In 1980, the OECD suggested that similar privacy guidelines[48] should apply to everyone handling personal information in transborder commerce. They updated these guidelines in 2013 in light of "changing technologies, markets and user behavior, and the growing importance of digital identities."[49] Interestingly, the OECD offers neither a definition of privacy, nor an argument about what constitutes a harm. What they offer are guidelines, compliance with which presumptively will prevent harm. The guidelines include these requirements (among other things):

- *There should be limits to the collection of personal data and any such data should be obtained by lawful means and, where appropriate, with the knowledge or consent of the data subject.*
- *Personal data should be relevant to the purposes for which they are to be used...*
- *The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfillment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.*
- *Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with [the previous requirement] except:*
  a) with the consent of the data subject; or
  b) by the authority of law.

The document states: "These Guidelines apply to personal data, whether in the public or private sectors, which, because of the manner in which they are processed, or because of their nature or the context in which they are used, pose a risk to privacy and individual liberties." The nature of the risk to privacy and liberty is left unstated–implicitly equated to inappropriate data disclosure.

The pragmatic utility of such guidelines is that, especially if translated into law or regulation, they create a context in which their violation becomes the harm in itself, a proxy harm that avoids the need to demonstrate that a specific harm to a data subject has occurred. Many have begun to point out the fundamental flaws with such consent-driven guidelines.[50] First, they embody the unstated assumption that

47. Bellovin, *Comments on Privacy.*
48. OECD, *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, `https://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm`, 1980.
49. OECD, *The OECD Privacy Framework*, `https://www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf`, 2013.
50. Bellovin, *Comments on Privacy.*

if the scope of use is specified, and the user has knowledge or has consented to that use, then no harm occurs. This assumption is difficult to defend in today's world. Notice and consent ignores in practice the large power of the data collectors, the information asymmetries that exist, their ability to pose "consent or go away" terms, and the downside to the data subject of failure to consent, especially when the data gatherer provides a highly useful application. [51] The user may not fully understand the future implications of consent, and what impact it might have. Second, this sort of blanket limitation may prohibit benign and innovative uses of data.

In 2011 the Digital Advertising Alliance proposed more aggressive guidelines with respect to harms, identifying specific uses that are explicitly prohibited: informing eligibility for employment credit, health care treatment, or insurance.[52] The guidelines recognize such uses as exploitations of the asymmetric relationship between an individual and a powerful service provider. The harm that could result is consistent with our definition of harm: an impairment to the welfare interests of, in this case, an individual. As we observed at the beginning of this topic, in the absence of codes of this sort, the holders of this information might feel compelled by competitive forces to exploit information in ways that are well outside the scope of delivering targeted ads, and in ways that have substantial potential to induce harm. The prohibitions reflect the fact that the user of services supported by targeted advertising is making an implicit bargain that his or her data can be gathered, but for limited purposes.

But how was this list generated? Is it complete? Who should have the responsibility to make such a list? Echoing our concerns about the abstractness of privacy as a principle, it does not seem possible to derive such a list top-down starting from a general privacy principle such as "the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others." That list emerges bottom-up, based on real-world examples of harmful behavior, and without use of the concept of privacy as a step in the reasoning.

We see two challenges in defining harms in this area is that a bottom-up list of harms can grow over time (as with case law) but will always be incomplete, and the use of consent is both amenable to abuse and may preclude benign and beneficial innovative uses of data. The lens of harm can help evaluate the acceptability (or unacceptability) of specific data use cases, but it leaves open the task of finding a legal basis to sustain those determinations. Privacy, including its purported tension with other considerations such as the First Amendment, provides one such basis,[53] but other factors are also relevant, such as balancing power between large corporations and individuals. We consider this issue in §5.

4.1 Targeted advertising

The current importance of advertising to the Internet ecosystem, as well as its impact on confidentiality, has brought together three issues—targeted advertising, data collection, and privacy—that might otherwise not be so closely linked. We consider each, and the threads that link them together, through the lens of harm.

The term *targeted advertising* refers to the delivery of advertisements based on knowledge about the recipient(s) – their interests, levels of wealth, past purchasing history, etc., today more or less without limit. Targeted advertising is not new (ads in mountain biking magazines differ from those in fishing magazines), and benefits markets, since it increases the probability that a presented ad is relevant and interesting to its viewer. However, the inevitable force to improve the level of relevance and interest of targeted ads creates the incentive to gather as much data as possible about the recipient, leading to a "creepiness" factor, which manifests in various ways, e.g., when ads follow a user from one site to another on the net, or as ads are delivered that seem ill-targeted. How the balancing of harms with free speech will play out in the context of targeted online advertising is not yet clear.

Perhaps the most compelling concrete fear of harms from data collection is that the data about individuals, once collected, will be used for other purposes. Absent limitations imposed by regulation or codes of conduct,

51. Feld addresses this point in his *cost of exclusion* metric of market power described in Harold Feld, *The Case for the Digital Platform Act: Market Structure and Regulation of Digital Platforms* (Roosevelt Institute / Public Knowledge, May 2019).

52. Digital Advertising Alliance, "Self-Regulatory Principles for Multi-Site Data," `https://digitaladvertisingalliance.org/sites/aboutads/files/DAA_files/Multi-Site-Data-Principles.pdf`, 2011,

53. Neil M. Richards, "Reconciling Data Privacy and the First Amendment," *UCLA Law Review* 52 (2005).

the incentive, if not the imperative, of the holder of such information would be to exploit it in whatever ways benefit the holder. The fear of undisclosed exploitation is informed repeatedly by discoveries of how large data collection companies behave. An example of exploitation of data collection that received much recent attention was the Facebook Cambridge Analytica scandal, where Cambridge Analytica obtained (via a researcher at the University of Cambridge, in violation of the university's terms of service with Facebook), profiles on 87 million Facebook users, which they used to influence voter sentiment in the 2016 presidential election.

Historical note on regulation of harms due to advertising

Advertising facilitates the effective functioning of markets, but some aspects of advertising are deemed harmful. However, in the United States, the First Amendment limits the government's ability to regulate speech. Rulings by the court provide some insights into circumstances where regulation of advertising (or more properly, *commercial speech*) is acceptable. The U.S. Supreme Court, in its Central Hudson ruling,[54] overturned a prohibition on advertising by electric utilities. In this case, the Court established a four-step analysis to determine the whether a ban on promotional advertising violated the First and Fourteenth Amendment:[55]

1. Is the expression protected by the First Amendment? For speech to come within that provision, it must concern lawful activity and not be misleading.
2. Is the asserted governmental interest substantial?
3. Does the regulation directly advance the governmental interest asserted?
4. Is the regulation more extensive than is necessary to serve that interest? (There must be a "reasonable fit" between the government's ends and the means for achieving those ends.)

In other words, since the purpose of commercial speech is informational, there can be no constitutional objection to the suppression of commercial messages that do not accurately inform the public. There is also no protection of commercial speech related to illegal activity. Outside these considerations, the State must assert a substantial interest achieved by any restriction on (even commercial) speech, and the regulation must achieve the goal in the most limited way possible. The FTC's authority to prohibit unfair or deceptive advertising reflects an understanding that those forms of advertising are harmful, and regulation to prevent them is acceptable. However, the *substantial interest* threshold reflects the reality that a determination of harm must reflect a balanced judgement, not the imposition of a simple bright line rule.

As another example of this balance, relying on the assumption that public health, and in particular the protection of children, is a matter where the U.S. government has a substantial interest, the FDA regulates and limits how tobacco and alcohol products can be advertised.[56] The U.S. government (the FCC) has also deemed harmful and thus imposed limits on the sheer volume of advertisements in children's television programming.[57] The FCC also prohibits TV stations from increasing the audio volume of commercial, relative to the shows they accompany.[58] In Kovacs v. Cooper, the Supreme Court ruled that regulations limiting or prohibiting 'loud or raucous' sound trucks on streets were constitutional, even though the sound trucks were projecting speech, including potentially advertising.[59] Sound trucks produce obnoxious levels of noise that anyone in proximity cannot avoid. So while the right to speech is protected, a speaker cannot take excessive steps to force people to listen. Note here it is not the advertising being regulated, but the harm induced by it.

4.2 Measurements and remedies: challenges and opportunities.

As we have noted, the abstract concept of privacy does not lead directly to the articulation of more specific harms. Thus, any attempt to measure *privacy harms* to an individual or group must be constructed bottom up, based on a particular class of incident. The concept of privacy seems more relevant to the construction of a remedy, rather than the claim of a harm. However, the concept of privacy is not the only legal tool available for harms that we associate with loss of privacy. Westin's definition of privacy included the claims of individuals, groups, or institutions. When private email of an organization is stolen, the harm is to the institution as a whole, as well as to the individual members. We note that in the U.S., which has no general privacy regulation, a remedy to such a harm could be pursued through other laws, such as the Computer Fraud and Abuse Act. Our emphasis on legal remedies reflects the structural character of the harms we covered in this section, which focused on established relationships between large data-driven companies and consumers, governed by notice and consent regimes. The measurement (i.e., investigatory) challenge is tracing how data is being used, and whether those uses are consistent with the consent given. Cases where there is blatantly inadequate notice and consent are perhaps easier to identify, compared to a violation of terms of service, such as the Cambridge Analytica example. These latter type of abuses are revealed by Reporters and researchers investigating specific circumstances. It is not clear how technical measurement can reveal these sorts of activities, but if other means can reveal them, the legal system may be able to identify the actor that should be held accountable.

5 Harms to Innovation, competition and choice

In this section we consider three interrelated aspirations: innovation, competition and choice, and economic growth, and harms that erect barriers to these goals. One reality with which we must contend is that despite the fascination with innovation as a driver of the economy, not all innovation is pro-consumer, and not all innovation leads to economic growth. Some innovations may lead to harms that justify regulation.

As a key component of the IT space, the Internet has contributed to economic growth by promoting innovation and creativity, technology development, and revolutionizing logistics and service industries, among other ecosystem disruptions. Specific properties of the architecture that facilitated these contributions include stability and openness of the architecture, most notably the narrow waist design that enabled innovation and experimentation below and above the waist, and specifications (open standards) that were (are) free from intellectual property restrictions. Innovation has been such a long-standing policy aspiration, indeed recognized as key to sustaining economic growth (despite how to accurately measure both being increasingly non-obvious), that it is generally elevated to a policy goal itself, and thus barriers to innovation recognized as at least ostensible harms. The question we address in this section is: what is distinctive about the digital era generally, or the Internet ecosystem more specifically, that raises new concerns about barriers to innovation?

To clarify, intrinsic barriers to innovation are not harms – innovation is exactly the overcoming of barriers. In the digital context, as in the past, a primary harm to innovation is anticompetitive behavior that blocks market entry. In current U.S. antitrust doctrine, a firm that fairly earns a dominant market position (or actual monopoly status) is not per se in violation of the law. In other words, the law does not consider that there is a harm to innovation from the failure of any specific business, entrepreneurial aspiration, or even entire industry faced with competitive disadvantage. But when a new company fails due to anticompetitive behavior of an incumbent rival, we say that the startup, and its stakeholders, have been harmed. Preventing and correcting the harms of anticompetitive behavior is the responsibility of antitrust law and enforcement.

54. Central Hudson Gas & Elec. v. Public Svc. Comm'n, 447 U.S. 557 (1980).
55. Wikipedia, *Central Hudson Gase and Electric Corporation vs Public Service Commission*, `https://en.wikipedia.org/wiki/Central_Hudson_Gas_\%26_Electric_Corp._v._Public_Service_Commission`, 2019, accessed July 25, 2019
56. For example, https://www.fda.gov/tobacco-products/products-guidance-regulations/advertising-and-promotion.
57. https://www.fcc.gov/consumers/guides/childrens-educational-television.
58. https://www.fcc.gov/consumers/guides/loud-commercials-tv
59. Kovacs v. Cooper, 336 U.S. 77 (1949)

Innovation is not limited to small business venturing and new market entrants. Innovation also occurs in large firms. However, for firms in dominant positions, the drivers of innovation may less compelling. Simplistically, motivations for innovation include growth, distinguishing a firm from its competitors, meeting changing needs (to keep existing customers), attracting the best talent, and reducing cost. Using those considerations, a firm that has achieved a dominant share of the market can no longer use innovation to substantially grow its market share. Innovation may allow a firm to branch into new services (which can generate new revenues from their existing customers), perhaps by buying smaller firms that have already developed an innovative offering. Innovation can reduce cost and improve business processes. But two key drivers of innovation are weakened for a dominant firm: the need to constantly fight competitors, and the benefit of growth in market share. Thus, the return on investment for innovation may be poor. This relationship reveals the natural linkage between healthy competition and vigorous innovation, which ties together the aspirations of innovation, competition, and choice among products and services.

The Four Principles of the FCC

The FCC issued a policy statement in 2005 in which it adopted these four principles:[60]

> [T]o ensure that broadband networks are widely deployed, open, affordable, and accessible to all consumers, the Commission adopts the following principles:
>
> - To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet, consumers are entitled to *access the lawful Internet content of their choice.*
> - To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet, consumers are entitled to *run applications and use services of their choice, subject to the needs of law enforcement.*
> - To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet, consumers are entitled to *connect their choice of legal devices that do not harm the network.*
> - To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet, consumers are entitled to *competition among network providers, application and service providers, and content providers.*

Freedom of choice also seems central to U.S. policy thinking – the term *choice* appears in three of the four FCC principles (see box), and *competition*, e.g., choice among providers, is in the fourth principle. But the word *choice* is ill-defined; people often invoke it as a proxy for some other goal such as consumer welfare, for which choice is either a means or a consequence. The logic is that competition leads to choice, and consumers will choose wisely, so competition disciplines providers toward offering products and services that consumers prefer. This line of reasoning implies that a loss of choice, measured as insufficient competition, is a harm, both to the economy generally and to the consumers who have fewer options among which to choose. However, the landscape of choice and competition is complex, and there are no straightforward criteria to judge when harm has occurred.

Some technical limitations of the Internet preclude certain areas of innovation. Until this decade, innovations in streaming video to the home were blocked by residential broadband access speeds. Some classes of applications do not work well on the public Internet—most problematic are those that require high reliability and availability, e.g., remote surgery, or remote control of autonomous vehicles. Internet enthusiasts stress its generality and its ability to support anything, and quietly ignore the applications it cannot so easily support. Society could accept that these applications are not going to happen, governments could attempt

60. Commission, *FCC 05-151, Policy Statement*

to incentivize changes in industry structure and coordination required to support them, or innovators might develop alternative, private networks that they could then exploit. This outcome would lead to competition between two industrial structures: the current Internet and this new alternative. If these applications were sufficiently compelling, more of what we see on the Internet today might migrate onto this new platform. This would be an example, at the lower packet-carriage layer of the Internet architecture, of a movement from a more decentralized and distributed structure to a network architecture with a more centralized character, controlled by a smaller number of more tightly coordinated (and likely powerful) actors.

5.1 The layered platform context for assessing harms to innovation

While the Internet has been a driver of many types of innovation, technology and economic forces in the Internet ecosystem tend to hinder choice and competition, and thus may slow innovation. In the era of plain old telephone service, the U.S. government considered provision of residential connectivity a natural monopoly. The government specifically encouraged AT&T, though its Bell Labs, to invest in innovation, which led to many wonderful inventions including the transistor. Today, government policy is to encourage competition in residential access, but in many locations – especially rural areas with low population density – competition in facilities-based access may not be economically viable, and monopoly may result, without the matching stimulus to innovate.

At the higher layers, economic forces also drive concentration and diminish choice. Once digital infrastructure is built, the marginal economics are dramatically different from those that have prevailed since the Industrial Revolution. Because fixed costs play such an important role in digital markets, they are characterized by disproportionately large returns to scale. Further, many digital markets are characterized by network effects that strengthen incumbents and weaken entrants: easy growth in scale, and easy expansion to reach global proportions. The result is a tendency of the market to tip in favor of one winner, leading to many well-known problems, including higher prices, less innovation, and lower quality in all its forms.

In the most recent decade, the proliferation of global *platforms* has driven Internet growth and use. The layered nature of the Internet ecosystem adds a nuance to the analysis of harms of concentration. Many (most) people have become sufficiently dependent on these platforms – personally and professionally – to consider them analogous to public utilities.[61] But the platform structure of the ecosystem, and these large companies that leverage this structure, creates options for anti-competitive behavior that differ from classic examples of harm, and both the platform structure and the information-centric conception of power suggest different harms and remedies. In a layered platform ecosystem, competition can occur among competing platforms, and it can also occur on top of the platform, where the complementors utilize the platform. Owners of dominant platforms may be able to exploit that platform to extract value from the complementors and shape the competition that occurs on top of the platform.

Harms from commercial API practices, and potential remedies

Mozilla's Chris Riley provides a useful illustration in the context of mechanism, process, or authority to challenge harmful practices regarding APIs:[62]

> APIs are the fundamental connective tissue of the internet. They are also a powerful tool for efficient, rapid scaling market entry, when a new app or service developer can reach users through existing APIs offered by platforms that have already achieved significant economies of scale. Yet, platform operators that have already hit a critical mass (and are thus less dependent for network effects on interconnection with others) face natural incentives to restrict the use of APIs by third parties. Some of these incentives are anti-competitive

61. The OECD recently profiled twelve of the world's leading platform companies to gain insight on what they do, how they do it, and why they succeed financially OECD, "An Introduction to Online Platforms and Their Role in the Digital Transformation," May 2019,

in intent and effect, for example if a platform operator obstructs a downstream market of services to its own detriment in order to prevent the growth of an emergent competitor. Others are driven by privacy and security concerns, for example shutting down third-party access to user data via an API rather than investing resources to determine how best to design the API and its policies and access controls to facilitate effective interconnection while also protecting privacy and security (and undertaking some risk of getting that balance wrong).

Many companies are already scaling back their API offerings. Facebook, most notably, has made major changes in the wake of the Cambridge Analytica scandal. Some of these changes, such as Facebook's deprecation of "publish_actions", have had significant and detrimental impact for smaller, independent technology projects.

One outcome of such rational decisions on the part of competing but dominant players, Riley imagines a plausible worst-case future scenario for the internet, where users choose from among a few non-interoperable vertically integrated technology stacks. The only recognizable "market" would be "internet services", and if five companies had roughly 20% market share each, it might qualify as a competitive market, but it certainly wouldn't be the internet as understood by people today. Riley emphasizes that it's not merely the size or multi-sided nature of companies that distinguishes the Internet from other industries. The most challenging unique feature in the context of competition is the "nature of the vertical integration of distinct and interconnecting digital services, and the fine-grained ability to control that interconnectivity through product and business decisions around integrating code bases and offering APIs — and how the outcomes those decisions can produce run counter to long-standing assumptions of interoperability and openness on which the internet was built." Riley also notes "Even before the term *platform* came into common parlance, that was how tech was designed — not in the two-sided market economic sense, but from the technical perspective that software and services are often built on top of other software and services built by others, relying on well-settled norms of openness and the mutual benefits of interoperability. Unfortunately, those norms are no longer settled, nor the mutual benefits guaranteed, in the digital economy prisoner's dilemma we have today."

Riley emphasizes the role of regulatory tools such as requiring the opening of private APIs to third parties, e.g., after a merger where two big companies make private APIs for each other. He considers interoperability requirements generally a more powerful tool than data portability requirements, given the network effects of data aggregators that cause huge barriers to entry.

As a special case of the platform ecosystem, the concept of the *multi-sided platform* is much in vogue today. A multi-sided platform is characterized by the fact that it supports multiple classes of interdependent users. It is the interdependency that makes the platform multi-sided. The platform operator can balance how it extracts value from the various classes of users, with the goal of extracting the maximum overall value. A dominant position in the market will enhance a platform's ability to extract such value. For example, in the case of the Apple app store, there is no way for application providers and users to interact except through the app store. Apple has prohibited the creation of competitive app stores, which gives Apple the power to exploit its platform to its advantage. In particular, Apple charges fees – in this case, to application developers – that some consider excessive. The Amazon store is a also a platform, one that sells Amazon-branded products that compete directly with third-party vendor products. These products frequently have preferred placement on the selling platform. Is this a harm that warrants intervention? How would one measure it?

There is thus a conundrum about the Internet: although the Internet itself was conceived as a general purpose platform to support any sort of application and service – in principle providing lots of choice – the actual economics of the ecosystem tip toward concentration. All competitive markets encounter the force of consolidation to some extent– the phrase "too big to fail" did not originate in the Internet industry. But it

62. Chris Riley, *Filing before the Federal Trade Commission*, March 2018, `https://blog.mozilla.org/netpolicy/files/2018/08/Mozilla-FTC-filing-8-20-2018.pdf`

bears asking, given the Internet's established status as critical infrastructure, how to minimize the effect of the failure of any one organization (or government?) on the operation of the internet? This returns us to the issue of resilience and risk assessment covered in §2.

Exploitation of the platform may rise to the level that constitutes a harm. Some examples may directly affect consumers, as in the case of the fee charged by Apple to app developers, which generally increase retail prices. But in many cases the consequence will be indirect and hard to measure. Apple has blocked certain applications from their app store, apparently for business reasons.[63] Is the harm to users that arises from their inability to get to one or another specific application on their Apple device so great as to justify intervention? In general, consumers understand that a given retail store cannot (and has no obligation to) carry all brands of product. Is the app store different in material ways, either because of the market power of Apple or because Apple is positioned in the digital ecosystem where inventory costs essentially nothing to create and essentially nothing to keep in stock?

5.2 Measurements and remedies: challenges and opportunities.

We consider three measurement challenges in the context of innovation: innovation itself, market power, and economic growth.

5.2.1 Measurement and remedies of harms to innovation

Measuring innovation (or lack thereof) would require agreement on proxy metrics, such as venture capital investments or IPOs, which are only a rough proxy for successful innovation. We consider the deeper and more vexing policy question in the context of harms – the societal benefit of different types of innovation. Two obvious paths to innovation for targeted advertising models are to learn more about those users so the platform operator can sell information about those users at a higher price, and to make those users "sticky"– lead them to spend more time on the platform so that they can be sold more often. As we discussed in §4.1, targeted advertising per se is not harmful—it is the potential for the misuse of the data collected that is harmful. But innovation that makes the user experience sticky by means of things like fake news and sensational content or even addictive games may be harmful not only to the user, but to society. Many aspects of innovation in this context may not be pro-consumer, nor in the best interests of society, although society may only be comfortable regulating exposure to such harms for children.

5.2.2 Measurement of market power and remedies of harms to competition

Measuring market power is complex in the context of multi-sided platforms. Do we measure their market share on the user-facing side because policy (the public interest) is primarily concerned with how the individual is affected by the firm's behavior? Do we just take the max of the share on the various sides? Can we consider each side separately, or do we enter into some complex multi-sided analysis, such as the analysis that occurred (in a different context) in Ohio vs. American Express?[64] With respect to competitive forces for innovation, in a multi-sided market, an actor may have a large market share on one side, but face competition on the other side. Consider Facebook, which dominates the social-media market on its user-facing side, but fights for advertising dollars on the other side. That is, although Facebook may dominate the social media market, advertisers can reach consumers by other channels. In the U.S. in 2018, Facebook captured 19% of digital advertising, while Google captured 36%.[65] We would expect to see innovation emerge on the side of their platform that faces competition, i.e., offers them the opportunity to capture a larger share of the advertiser market. Indeed there is innovation in this space, which takes the form of better grooming their users to improve Facebook's ability to sell these users to the advertisers.

63. Apple's curation of the app store, refusing to include apps that seem to have harmful consequences, would seem a benign form of blocking that improves the user's level of trust.

64. Ohio v. American Express Co., 585 U.S. ___ (2018)

65. eMarketer, *Data Suggests Surprising Shift: Duopoly Not All-Powerful*, March 2018, `https://www.emarketer.com/content/google-and-facebook-s-digital-dominance-fading-as-rivals-share-grows`.

More generally, should we conceptualize market power differently in the information age? Is there a form of market power that relates to the share of information rather than share of market? Colloquially, it makes sense that a firm that has control of a large quantity of information is in a powerful position. But is there a way (or should we look for a way) to map that into some concept of market power or dominance? If there were such a mapping, might it happen that a firm that does not have dominance in terms of market share could still be dominant with respect to information? Are these independent measures or inter-related?

Some remedies proposed to deal with market power in the information age reflect dominance of information. These remedies include requiring firms with large quantities of information to share it with rivals, perhaps for a fee, or perhaps as an alternative to paying taxes. Firms that hold personal information might be required to release that information to the person to whom it applies. Even if these ideas are speculative and would require thought to reduce to practice, they suggest that the information age will require reconception of remedies to market power.

As an example of such reconception, Feld proposed a new metric of market power known as the *cost of exclusion*, or CoE.[66] The focus is on the cost to a third party firm or user for being excluded from the platform. This measure avoids any necessity of talking about market share, or whether the platform is multi-sided. The measure is specifically relevant to the layered platform structure of the Internet ecosystem, and more directly captures the aspect of market power that can lead to harms.

Harms that arise from market power traditionally fall under the jurisdiction of antitrust enforcement. In the context of layered platforms in the digital ecosystem, forms of anticompetitive behavior may differ, so the justification for intervention and potential remedies may need reconsideration. The advertising-supported ecosystem illustrates the challenge of justifying intervention under current antitrust doctrine, which tends to anchor on consumer-facing prices and output. When the price for a service (like Facebook) in dollars is zero, market power does not lead to higher nominal prices. What occurs instead is likely more accurately characterized as a degradation in the *quality-adjusted* price, where innovations that lead to addiction and stickiness and more precise tracking and profiling represent losses of quality. If a platforms price is zero but the quality of the service improves, then its quality-adjusted price falls. Conversely, it the quality falls, the quality-adjusted price rises.

However, it seems very hard to quantify those losses. It might be possible to quantify the gains to the platform (since these innovations map in some way to revenues), but the loss to the consumer is hard to estimate, which raises a basic question regarding net benefit or welfare generated, looking across both sides of the platform. Antitrust doctrine in the U.S. for the last several decades has tended to limit its analysis to objective measures such as changes in short-run price and output, which has minimized scrutiny of dominant two-sided platform companies offering zero-priced services (where "you pay with your data"). Many have begun to analyze how market definitions, user-facing prices, the role of data, and the benefits of innovation challenge competition regulators who are responsible for evaluating mergers and conduct in the tech sector through the lens of advancing consumer welfare. Several have called for expansion of the regulatory toolset.[67] Caves and Singer have recently argued the merits of addressing innovation harms outside of antitrust doctrine, using a non-discrimination standard. They found a reasonable basis to conclude that by requiring concrete evidence of consumer injury attributable to the exclusionary conduct, the consumer welfare standard leads to under-enforcement in the area of innovation harms as well. Indeed, they note that the federal agencies have not a pursued a pure innovation-based theory of harm under Section 2 of the Sherman Act in nearly two decades.

Others have described more fundamental harms to self-determination and autonomy in the future.[68]

66. Feld, *The Case for the Digital Platform Act: Market Structure and Regulation of Digital Platforms.*

67. Bruno Lasserre and Andreas Mundt, "Competition Law and Big Data: The Enforcers' View," 2017, `%5Chttps://www.bundeskartellamt.de/SharedDocs/Publikation/EN/Fachartikel/Competition_Law_and_Big_Data_The_enforcers_view.pdf`; Kevin Caves and Hal Singer, "When the Econometrician Shrugged: Identifying and Plugging Gaps in the Consumer Welfare Standards," *George Mason Law Review* 26, no. 2 (2018).

68. Shoshana Zuboff, *The Age of Surveillance Capitalism* (2019).

5.2.3 Measurement of economic growth and remedies of harms

Although the potential for increased productivity from information and communication technologies seem enormous, actual statistics about productivity growth are far less sanguine. The average annual growth in productivity per hour dropped from 2.8 percent in the mid-twentieth century to 1.6 percent as the digital era began.[69] In exploring root causes of this drop, Gordon observes that advances since 1970 have been narrowly distributed into sectors such as entertainment, communications, and data mining. As selling data-mined information about consumers' interests became the dominant economic model of the Web, consumer-facing, as opposed to productivity-enhancing, activities thrived. Innovations based more in atoms than bits, e.g., food, clothing, shelter, transportation, health, and working condidtions both inside and outside the home, fell behind.

The implication is that innovations to better capture advertising may not lead to overall economic growth. The pot of advertising dollars can indeed grow to some extent. Today, online advertising is over half of all advertising spend, so if the online ecosystem displaced all other forms of advertising it might grow by a factor of two. Beyond that, innovation in how firms like Facebook sell users to advertisers will not drive growth in advertising. In general, the advertising spend for any firm is some fraction of their total revenue.[70] Firms want to grow their top line while not growing expenses that drive down the bottom line. Firms that compete to attract advertising dollars are, to first order, fighting over a pot of money the size of which is controlled by other actors. It is, from the perspective of those firms and in the short run, a fixed size pot. Indeed, if online advertising leads to growth in consumer spending, then it contributes to economic growth, and to more advertising dollars. So there may be a long-run growth cycle, but it is not clear that the innovations we see in tracking and profiling lead to more consumer spending; they may just lead to more revenues per ad.

Furthermore, while the U.S. economy has been expanding for a decade, wage growth has been slow, particularly for less-skilled workers, while labor force participation has remained low. Krueger and Posner have attributed these macroeconomic effects to increased monopsonization of and/or collusion in labor markets, which in turn they attribute to a combination of (1) merger activity spanning several decades; and (2) the rise of industries with strong network effects, creating massive employers with power in labor markets; and (3) non-compete and non-poaching agreements.[71] They also found no strong relationship between concentration and average wages, suggesting that the increased productivity of superstar firms is is not passed through to labor in the form of higher compensation.

6 Harms to journalism, the marketplace of ideas, and the political processes that depend on them

Accusations against media are thriving in the Internet age: polarization and filter bubbles, and related lack of diversity in news sources and ideas, fake news, death of journalism, ironically accompanied by constant addiction to media feeds that lead to a new (some say degraded) social fabric of continuous partial attention. The ramifications of these difficulties lead to the most daunting harms yet – to our ability to engage in trustworthy democratic political processes. This harm is paramount because it undermines society's ability to develop just and effective remedies for not only this harm but any other harm.

Despite its obvious importance, we will only briefly cover this topic. It is well outside our expertise, and it has been recently thoroughly covered by others. Perhaps most notably to date, Harold Feld offers a comprehensive assessment and remedies for the landscape of media-related harms attributed at least in part to the rise of new digital media platforms in Chapters 5 and 6 of his recent book, *A Case for the Digital*

69. Robert Gordon, *The Rise and Fall of American Growth* (Princeton University Press, 2016).

70. See *The CMO survey: Results by Firm & Industry Characteristics*, 2018, pp 57-58, `https://cmosurvey.org/wp-content/uploads/sites/15/2018/08/The%5C_CMO%5C_Survey-Results%5C_by%5C_Firm%5C_and%5C_Industry%5C_Characteristics-Aug-2018.pdf` for data about advertising spend as a fraction of revenue for firms in different industry sectors.

71. "A Proposal for Protecting Low-Income Workers from Monopsony and Collusion — The Hamilton Project," accessed August 18, 2019, `https://www.hamiltonproject.org/papers/a_proposal_for_protecting_low_income_workers_from_monopsony_and_collusion`.

Platform Act.[72,73] His book draws on 100+ years of history of communications and media law, which reference harm from digital media platforms in the context of its impact on democracy – and the marketplace of ideas. The growing zeitgeist he captures is the danger to democracy from having consolidated media platform firms with so much power. It is a return of communications policy to communications theory – recognition of what these networks are really about – the exchange of information and the interaction of humans.

Wheeler, who wrote a forward to Feld's book, also reviewed highlights of the history of news media in his own book *From Gutenberg to Google*.[74] He observed that news media was rather local and biased until the advent of the first wide area electronic network, the telegraph. This infrastructure triggered the first commercialization of news gathering, fueled by advertising revenues. A simple reality emerged: the more readers, the higher the advertising revenue, which encouraged balanced reporting in order to offend as few readers as possible. As the Internet began to to swallow journalism, it became less costly to produce or distribute news, putting tremendous pressure on news media organizations, another force for consolidation. Klinenberg's 2007 book *Fighting for Air*[75] (also covered in Feld's blog series) describes the continued horizontal and vertical consolidation of newspapers, accompanied by undeniable metrics of market failure in terms of reporter layoffs and reductions in substantive analysis. The weak economics of the industry has been repeatedly used to justify further consolidation, for the last several decades.

Feld emphasizes that many other aspects of media history are repeating themselves, including questions as to what responsibilities the media bears to to maintaining a democratic society, how to protect consumers from potential harms (fraud, abuse, and gravitation toward monopoly). Each new media technology, from telegraph to broadcasting to telephony, experiences a transition from an initial "techno-euphoria" stage to a phase of concern as the industries concentrate toward enormous political and social power. But he cautions against the idea that new digital platforms and news aggregators should subsidize existing media conglomerates. As he notes, "Any effort to address the underlying crisis of solvency and the crisis of faith in modern news media must acknowledge the role that news conglomerates have played in creating these crises, not merely point fingers at platforms."

6.1 Measurement challenges and remedies

Feld proposes many remedies to the challenges of consolidating media, the first set of which are primarily about how to enhance or preserve competition in this industry, and thus overlap with our discussion of innovation (§5). He proposes a new regulatory body, which he calls the Digital Platform Agency, which would have at its disposal several tools, including structural remedies such as legislative limits on market share, product unbundling and structural separation requirements, and non-discrimination requirements. Higher up the platform layers, he also offers content-related remedies, to address diversity, censorship, and content moderation challenges.

1. Require an open API by social media platform algorithms, so that independent third parties could create public-interest algorithms to understand the effects of social media distribution patterns, similar to how public interest groups monitor the mainstream media today.

2. To remedy the risk of filter bubbles, select recommendations through "wobbly" algorithms that provide a wider range of possible results.

3. To promote representational diversity, prohibit algorithms from applying certain suspect classifications or actively reversing these categories at random. (E.g., at random intervals, the algorithm should assume the user is female rather than male while holding all other factors the same.);

4. Promote news and media literacy as components of basic education.

72. Feld, *The Case for the Digital Platform Act: Market Structure and Regulation of Digital Platforms*.

73. A useful preamble to this work was his three-part blog series in 2018, "We Need to Fix the News Media, Not Just Social Media" Harold Feld, *We Need To Fix The News Media, Not Just Social Media. A 3-part essay.*, `https://wetmachine.com/tales-of-the-sausage-factory/we-need-to-fix-the-news-media-not-just-social-media-part-i/`, 2018.

74. Tom Wheeler, *From Gutenberg to Google: The History of Our Future* (Brookings Institution Press, 2019).

75. Eric Klinenberg, *Fighting for Air: The Battle to Control America's Media* (Metropolitan Books, 2007).

5. Encourage the research and development of new tools to identify reliable sources of information.

His suggestions also reflect the idea that the actor in the best position to remedy a harm is not necessary that actor that caused the harm.

7 Conclusions

Our original motivation for this paper was pragmatic and operational. The background we bring to this paper is on the one hand architectural—how the Internet ecosystem is actually structured—and on the other hand empirical—how we should measure the Internet to best understand what is actually happening. If everything were wonderful about the Internet today, the need to measure and understand would not be so compelling. A justification for measurement follows from its ability to shed light on problems and challenges. Sustained measurement or compelled reporting of data, and the analysis of the collected data, generally comes at considerable effort and cost, so must be justified by an argument that it will shed light on something important.

This reasoning naturally motivates our taxonomy of things that are wrong–what we call *harms*. That is where we, the research community generally, and governments should focus attention. We do not intend this paper as a catalog of pessimism, but to help define an action agenda for the research community and for governments.

The structure of the paper proceeds along a path that moves up the conceptual layers of the Internet architecture and the political economy in which it is embedded, from technology to society. We draw primarily on our technology expertise – for harms that are closer to the technology, we can be more specific about the harms, and more specific about possible measurements, remedies, and actors that could undertake them. We also acknowledge broad societal or economic harms, which are neither strongly shaped by the specifics of the Internet architecture nor as amenable to study via technical measurement of the Internet itself, at least in its traditional forms. We wanted to be methodical in finding harms, to avoid missing important issues. But we assume that our list is not complete; this paper is a conversation in progress, and we hope it inspires feedback.

7.1 Barriers and incentives to improving empirical grounding for policy

Our emphasis on measurement comes from a belief that both good science and good policy depend on sound interpretation of good data. Good policy also depends on a careful analysis of what constitutes a harm, which draws on principles of philosophy and ethics, but also on common sense. Measurement plays multiple roles here. In some cases it can help determine if a harm is occurring, but it can also help inform the consideration of and remedies. Measurement can lend confidence that a given remedy is effective, acknowledging the inferential challenge that correlation is not causality. Measurement can also inform priority-setting of policymakers.

Various actors engage in Internet measurement today. ISPs measure aspects of their own networks, but treat what they learn as proprietary. Third parties, such as academic researchers, undertake measurements from vantage points external to ISP networks, albeit with limited visibility. Data related to many of the harms we have covered will require mandatory reporting by relevant actors to relevant agencies. In some cases, it may be possible to release data to to the research community under strict guidelines, similar to how researchers access raw census data. In other cases, governments or dedicated organizations under the auspices of governments may have to perform the data analysis.

Broadband coverage mapping is an example that has received particular attention in recent U.S. administrations, in part because the government allocates federal funds to connect unserved regions based on the data, and yet there are serious questions as to the accuracy of the coverage data. A more scientific approach to measuring availability-related metrics – e.g., curating and validating data from carriers, census data on households, sampling actual web speed (bandwidth) measurements – would require significant investment and coordination than Congress or the FCC have allocate thus far. We expect greater calls for such investment in the future, perhaps subsidized by industry segments with the most to gain from having more

accurate data. We also expect organizations like (in the U.S.) the Census Bureau, the National Bureau of Economic Research, and even the Bureau of Labor Statistics to develop operational interest in integration of Internet-ecosystem metrics into their current analysis and reporting.

Third-party measurement efforts are often opportunistic, taking advantage of some feature of the ecosystem to gather data, perhaps in ways not intended by the creator of the feature. Such efforts may put the researcher into conflict with the operator of the system, who may not want to reveal what the researcher is discovering, however useful the research results are. This problem is most acute in empirical security research, where researchers attempt to understand system vulnerabilities. Several laws (DMCA, CFAA, ECPA) put serious constraints on such research and analytics that would allow understanding and tracking of many security-related harms. The net effect is that in many cases security research is effectively criminalized, if that research requires or results in access to a network that has not explicitly authorized the access, even if the intent of the research is to benefit the security of that network and its users. To the extent that improved understanding is prerequisite to assessing the scope and extent of security harms, we also need remedies for these current barriers to scientific research. One proposed remedy has been that governments reshape laws to focus on the intent underlying the activity. That is, if an academic researcher can certify the intent of their activity is for research, and if it did not cause any demonstrable harm, it should not be illegal. Conversely, if online actions are intended to create or promote malicious activity, those actions should be illegal. The onus should be on the actor to prove intent.[76]

7.2 A review of harms and remedies

Our primary goal in this paper was to catalog harms, so as to inform and focus both policy-makers and the research community. We took a structural approach to our taxonomy—where in the ecosystem does the harm arise, and what actors bear responsibility for the harm, or are best positioned to mitigate the harm. We emphasized that is often the case that the actor that causes the harm is not the one that can best mitigate it. This is most obvious in the case of intentional attacks, where we do not expect the actor that causes the harm (the attacker) to play any role in mitigating it. More generally, and following the tradition of tort law, the actor that should be burdened with the task of mitigating the harm should be the one in the best position to do so. The burden of mitigation does not derive from any fault of that actor, but from the position of that actor in the ecosystem. So our structural approach to organizing harms, with a focus on where the harm occurs, helps to identify good strategies for mitigation because it can help identify which actors live in that part of the ecosystem, and are thus good candidates to burden with the task of mitigation. As in other ecosystems, the actor that is in a good position to mitigate a harm may have no incentive to do so; thus mitigation of certain harms will require regulation or other incentive for the relevant actor to undertake mitigation. In some cases, we also identified possible remedies—approaches to mitigating certain harms. We discussed remedies in the context of specific harms, but the relationship is not one-to-one. Many remedies target more than one harms, e.g., Feld's proposed use of CPNI regulations on digital platform companies to remedy issues with both competition and privacy.[77]

7.2.1 Harms to access

- Lack of universal service. The lack of any physical access in certain places or to certain people. The measurement challenge relates to accuracy of mapping. The general form of the remedy involves direct involvement and investment by the public sector to complement what the private sector can justify.

- Low quality service. Service is of inadequate quality to support the needs of a typical user. Independent measurements can gather some aspects of service quality. Competition in access service might stimulate the provider to make a better offering, but this harm may occur exactly where competition is not thriving. In the context of universal service, where the public sector is investing, careful specification

76. Andrew Burt and Daniel E Geer, "Flat Light: Data Protection for the Disoriented, from Policy to Practice," `https://www.hoover.org/sites/default/files/research/docs/burtgeer_flatlight_revisednov20_webreadypdf_final.pdf`, 2018,

77. Feld, *The Case for the Digital Platform Act: Market Structure and Regulation of Digital Platforms.*

of required service quality can prevent this harm. If the harm arises in the context of a privately-built network, the only remedy may be direct public sector investment to augment infrastructure capabilities.

- High cost service. If competition is lacking, the only remedy is either price regulation or direct public sector investment to cover costs.

- Low adoption. To some extent, governments can address this harm through education, and subsidies where cost is the barrier to adoption. The population of non-users may also naturally shrink over time.

- Insufficient resilience. This harm is more complex. Measurement is challenged by the distributed and decentralized nature of service provision in the ecosystem, although much consolidation has happened over the last two decades. Still, each service provider makes its own decisions about degree of redundancy, degree of interconnection with other regions, and so on. An ambitious measurement campaign, probably requiring the creation of a dedicated organization with governmental backing, may be able to learn something about the resilience of the current Internet by gathering massive amounts of routing data, but this would not paint a complete picture, since operators can physically reconfigure their networks in response to disasters. While governments can encourage, or require by regulation, operators to undertake a certain degree of redundancy in their parts of the network, application designers can also play a role by designing mechanisms in applications that compensate for failures at the lower layers.

7.2.2 Harms to integrity–loss of trust

- Physical and network layer. We identified three critical systems in the Internet that can lead to misrouted and misdelivered traffic: BGP, the DNS and the CA system. The state of all three of these systems could in principle be measured, but only with great effort, again probably requiring a dedicated organization with government backing, and the cooperation (voluntary or mandated) of many actors in the ecosystem. The decentralized character and lack of coordination hinders the viability of many proposed remedies to abuses of these systems. In the case of the CA system, Google has used its position in the market to impose a remedy that depends on more centralized control. End-nodes have only a limited ability to protect themselves from these harms.

- Edge devices and routers. The penetration of an edge device can lead to many consequential harms, including ransomware and other malware, also covered at the application layer. Measurement of data breach was not practical until a law required

  that holders of data disclose data breaches.

- Application layer. We identified three baskets of harms. The first, malicious applications, presents challenges to detection. It is often hard to tell what an application is doing, especially if encrypted. We also described applications that, while not themselves malicious, allow interactions among users without adequate protection against malicious users. There is no general method for assessment of such harms. The primary responsibility for mitigating these harms will have to fall on the designers of the applications, with nudging by industry groups or regulation by governments as deemed necessary.

7.2.3 Harms to confidentiality—privacy

Ironically, privacy is an abstract concept in today's personal-data-driven ecosystem, and loss of privacy is not generally amenable to measurement. Law and regulation has provided a proxy harm, which is violation of requirements for notice and consent. One dimension of measurement would be to assess the degree of clarity of different notices about the terms of data disclosure and use. Misuse and abuse of data are likely to be covert activities, discovered by investigators and journalists. As with other forms of crime, a clear taxonomy of the crimes, and clear reporting requirements, will be prerequisite to overall assessment of the level of harm.

7.2.4 Harms to innovation, competition, market power, and economic growth

- Innovation. Measurement of innovation uses economic proxies such as rate of sector-specific IPOs. The deeper question is the how different innovations may lead to benefit or harm, including in terms of economic growth. We propose some structural answers, but this area is largely outside our scope of measurement expertise, and we leave it to future conversations.
- Market power. Measurement of market power will require new approaches to accommodate the dynamics of "information capitalism", and since multi-sided platforms may have different degrees of market share on the different sides. Feld's recent book proposes and defends a new metric of market power, called Cost of Exclusion, which, as he describes it, maps more directly to the potential of harm from market power of digital platforms.

7.2.5 Harms to journalism, the marketplace of ideas, and the political processes

We acknowledged the existence of high-level societal harms, which arise not from the specifics of the Internet architecture, but from more general and fundamental characteristics of the digital ecosystem—such as scale-free expansion or the network effects that drive toward consolidation. We refer the reader to Feld's recent work to identify a number of potential remedies to mitigate some of harms that result from these tendencies.

7.3 A final thought

One motivation for this paper is that we believe the Internet ecosystem is at an inflection point. The Internet has revolutionized our ability to store, move, and process information, including information about people, and we are only at the beginning of understanding its impact on society and how to manage and mitigate harms resulting from unregulated commercial use of these capabilities. Current events would suggest that now is a point of transition from laissez-faire to regulation. However, the path to good regulation is not obvious, and now is the time for the research community to think hard about what advice to give the governments of the world, and what sort of data can back up that advice. Our highest-level goal for this paper is to contribute to a conversation along those lines.